



US008280905B2

(12) **United States Patent**
Shields et al.

(10) **Patent No.:** **US 8,280,905 B2**
(45) **Date of Patent:** **Oct. 2, 2012**

(54) **AUTOMATED FORENSIC DOCUMENT SIGNATURES**

(75) Inventors: **Thomas Clay Shields**, Washington, DC (US); **Ophir Frieder**, Chicago, IL (US); **Marcus A. Maloof**, Bethesda, MD (US)

(73) Assignee: **Georgetown University**, Washington, DC (US)

(*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 570 days.

(21) Appl. No.: **11/963,186**

(22) Filed: **Dec. 21, 2007**

(65) **Prior Publication Data**

US 2009/0164517 A1 Jun. 25, 2009

(51) **Int. Cl.**

G06F 7/00 (2006.01)
G06F 17/30 (2006.01)

(52) **U.S. Cl.** **707/781**; 706/46; 713/186; 713/180

(58) **Field of Classification Search** 706/46
See application file for complete search history.

(56) **References Cited**

U.S. PATENT DOCUMENTS

4,697,209	A	9/1987	Klewit et al.	
5,960,081	A	9/1999	Vynne et al.	
6,014,183	A	1/2000	Hoang	
6,078,917	A *	6/2000	Paulsen et al.	707/999.005
6,263,349	B1 *	7/2001	Anderson	1/1
6,279,010	B1 *	8/2001	Anderson	707/999.202
6,345,283	B1 *	2/2002	Anderson	707/999.001
6,683,966	B1	1/2004	Tian et al.	
6,772,196	B1 *	8/2004	Kirsch et al.	709/206
7,058,968	B2 *	6/2006	Rowland et al.	726/1
7,080,091	B2	7/2006	Matsuda	
7,200,244	B2	4/2007	Yacobi	

7,206,303	B2 *	4/2007	Karas et al.	370/350
7,260,722	B2	8/2007	Forstrom et al.	
7,289,643	B2	10/2007	Brunk et al.	
7,359,006	B1	4/2008	Xiang et al.	
7,506,379	B2 *	3/2009	Abali et al.	726/26
7,738,704	B2	6/2010	Lienhart et al.	
7,739,114	B1	6/2010	Chen et al.	
7,742,617	B2 *	6/2010	Smith et al.	382/100

(Continued)

FOREIGN PATENT DOCUMENTS

WO WO-2007/075813 7/2007

OTHER PUBLICATIONS

Hines, Matt, "Provilla to Fight Leaks with Document Fingerprints", www.eweek.com/article2/0,1985,2087164,00.asp, Ziff Davis Holdings Inc., May 12, 2008, 4 pgs.

(Continued)

Primary Examiner — James Trujillo

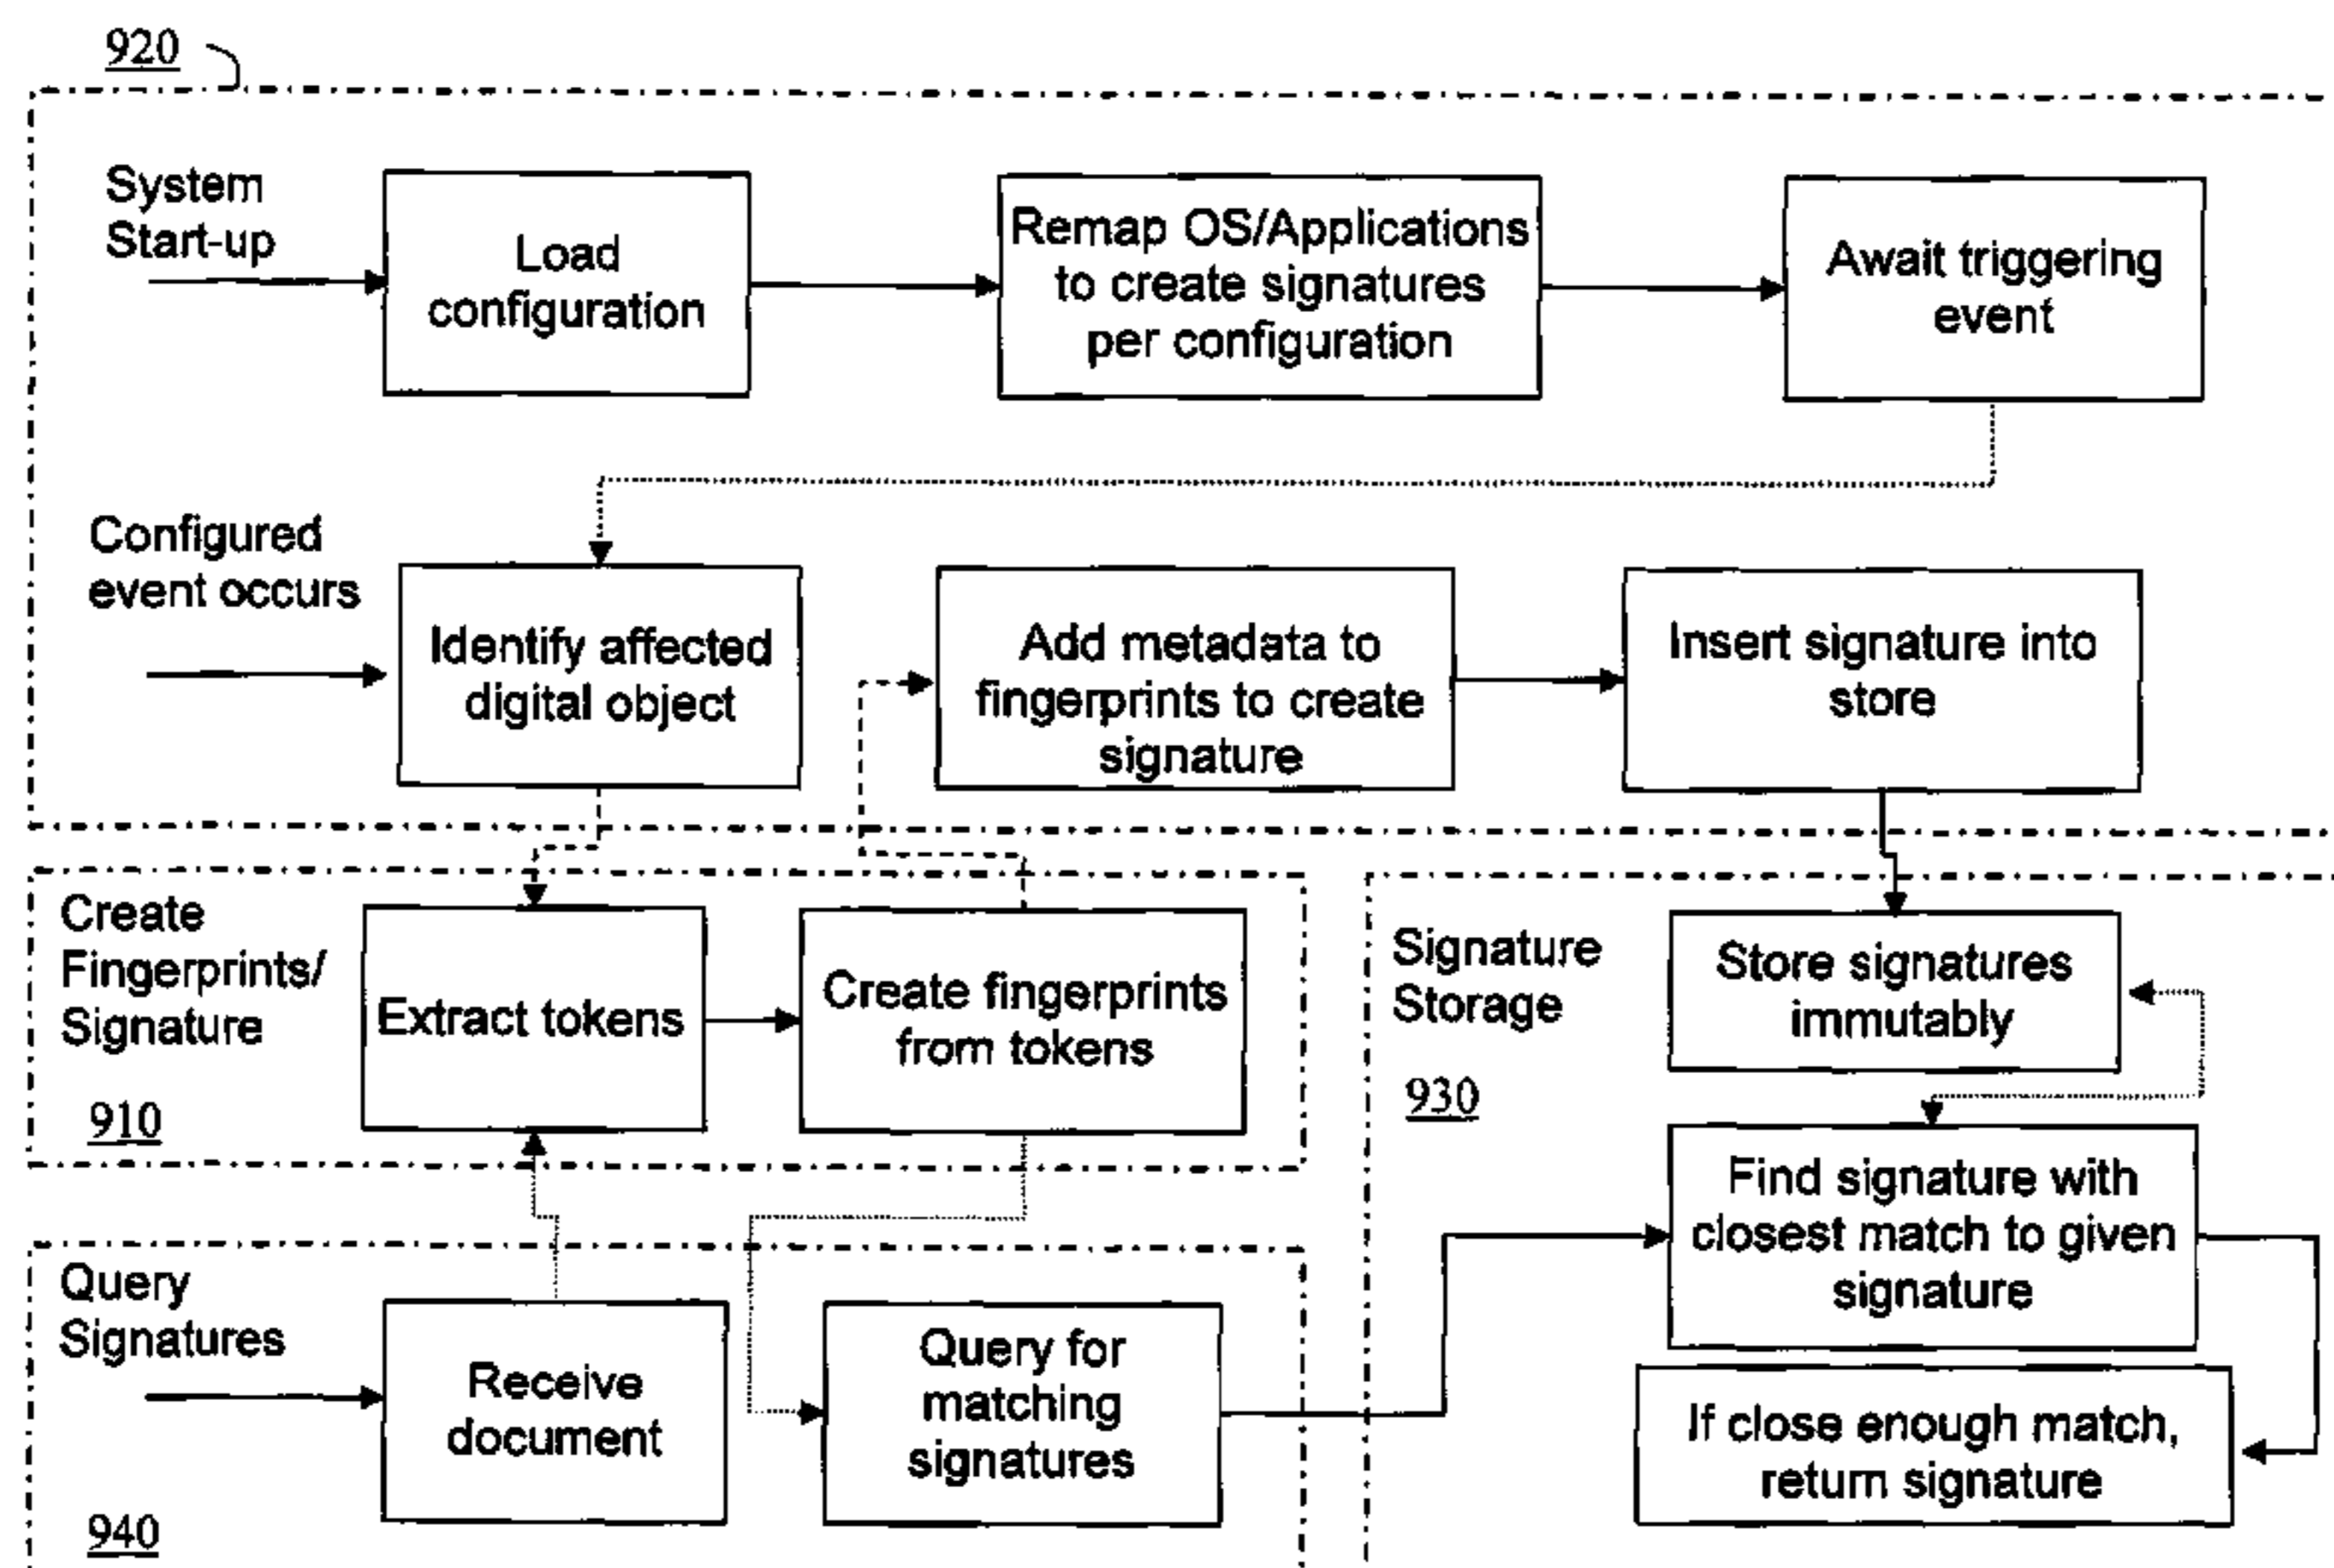
Assistant Examiner — Pavan Mamillapalli

(74) *Attorney, Agent, or Firm* — Matthew K. Ryan; Brian M. McGuire; Frommer Lawrence & Haug LLP

(57) **ABSTRACT**

Methods and systems are provided for a proactive approach for computer forensic investigations. The invention allows organizations anticipating the need for forensic analysis to prepare in advance. Digital signatures are generated proactively for a specified target. The digital signature represents a digest of the content of the target, and can be readily stored. Searching and comparing digital signatures allows quick and accurate identification of targets having identical or similar content. Computational and storage costs are expended in advance, which allow more efficient computer forensic investigations. The present invention can be applied to numerous applications, such as computer forensic evidence gathering, misuse detection, network intrusion detection, and unauthorized network traffic detection and prevention.

35 Claims, 9 Drawing Sheets



U.S. PATENT DOCUMENTS

7,899,205 B2* 3/2011 Pelly et al. 382/100
7,954,151 B1* 5/2011 Nisbet et al. 726/22
2002/0116392 A1 8/2002 McGrath et al.
2003/0103645 A1 6/2003 Levy et al.
2003/0105739 A1* 6/2003 Essafi et al. 707/1
2003/0110396 A1* 6/2003 Lewis et al. 713/201
2003/0149744 A1 8/2003 Bierre et al.
2004/0028281 A1 2/2004 Cheng et al.
2004/0133548 A1 7/2004 Fielding et al.
2004/0260733 A1* 12/2004 Adelstein et al. 707/202
2005/0193005 A1 9/2005 Gates et al.
2005/0273614 A1* 12/2005 Ahuja et al. 713/176
2006/0031938 A1* 2/2006 Choi 726/25
2006/0101069 A1 5/2006 Bell et al.
2006/0153296 A1 7/2006 Deng
2006/0253423 A1 11/2006 McLane et al.
2006/0253438 A1* 11/2006 Ren et al. 707/5
2006/0256739 A1* 11/2006 Seier et al. 370/261
2007/0074147 A1 3/2007 Wold

2007/0085710 A1* 4/2007 Bousquet et al. 341/50
2007/0085711 A1* 4/2007 Bousquet et al. 341/50
2007/0253594 A1 11/2007 Lu et al.
2007/0260643 A1 11/2007 Borden et al.
2007/0276823 A1* 11/2007 Borden et al. 707/5
2007/0283158 A1* 12/2007 Danseglio 713/180
2007/0294271 A1 12/2007 Bammi et al.
2008/0016569 A1* 1/2008 Hammer et al. 726/23
2008/0065811 A1* 3/2008 Jahangiri 711/100
2008/0163288 A1 7/2008 Ghosal et al.
2008/0184338 A2* 7/2008 McCreight et al. 726/3
2009/0216761 A1 8/2009 Raichelgauz et al.
2011/0047130 A1* 2/2011 Kim et al. 707/692

OTHER PUBLICATIONS

Brian Carrier; The Sleuth Kit Informer; Issue #6; Jul. 15, 2003; pp. 1-7.

* cited by examiner




FIG. 1

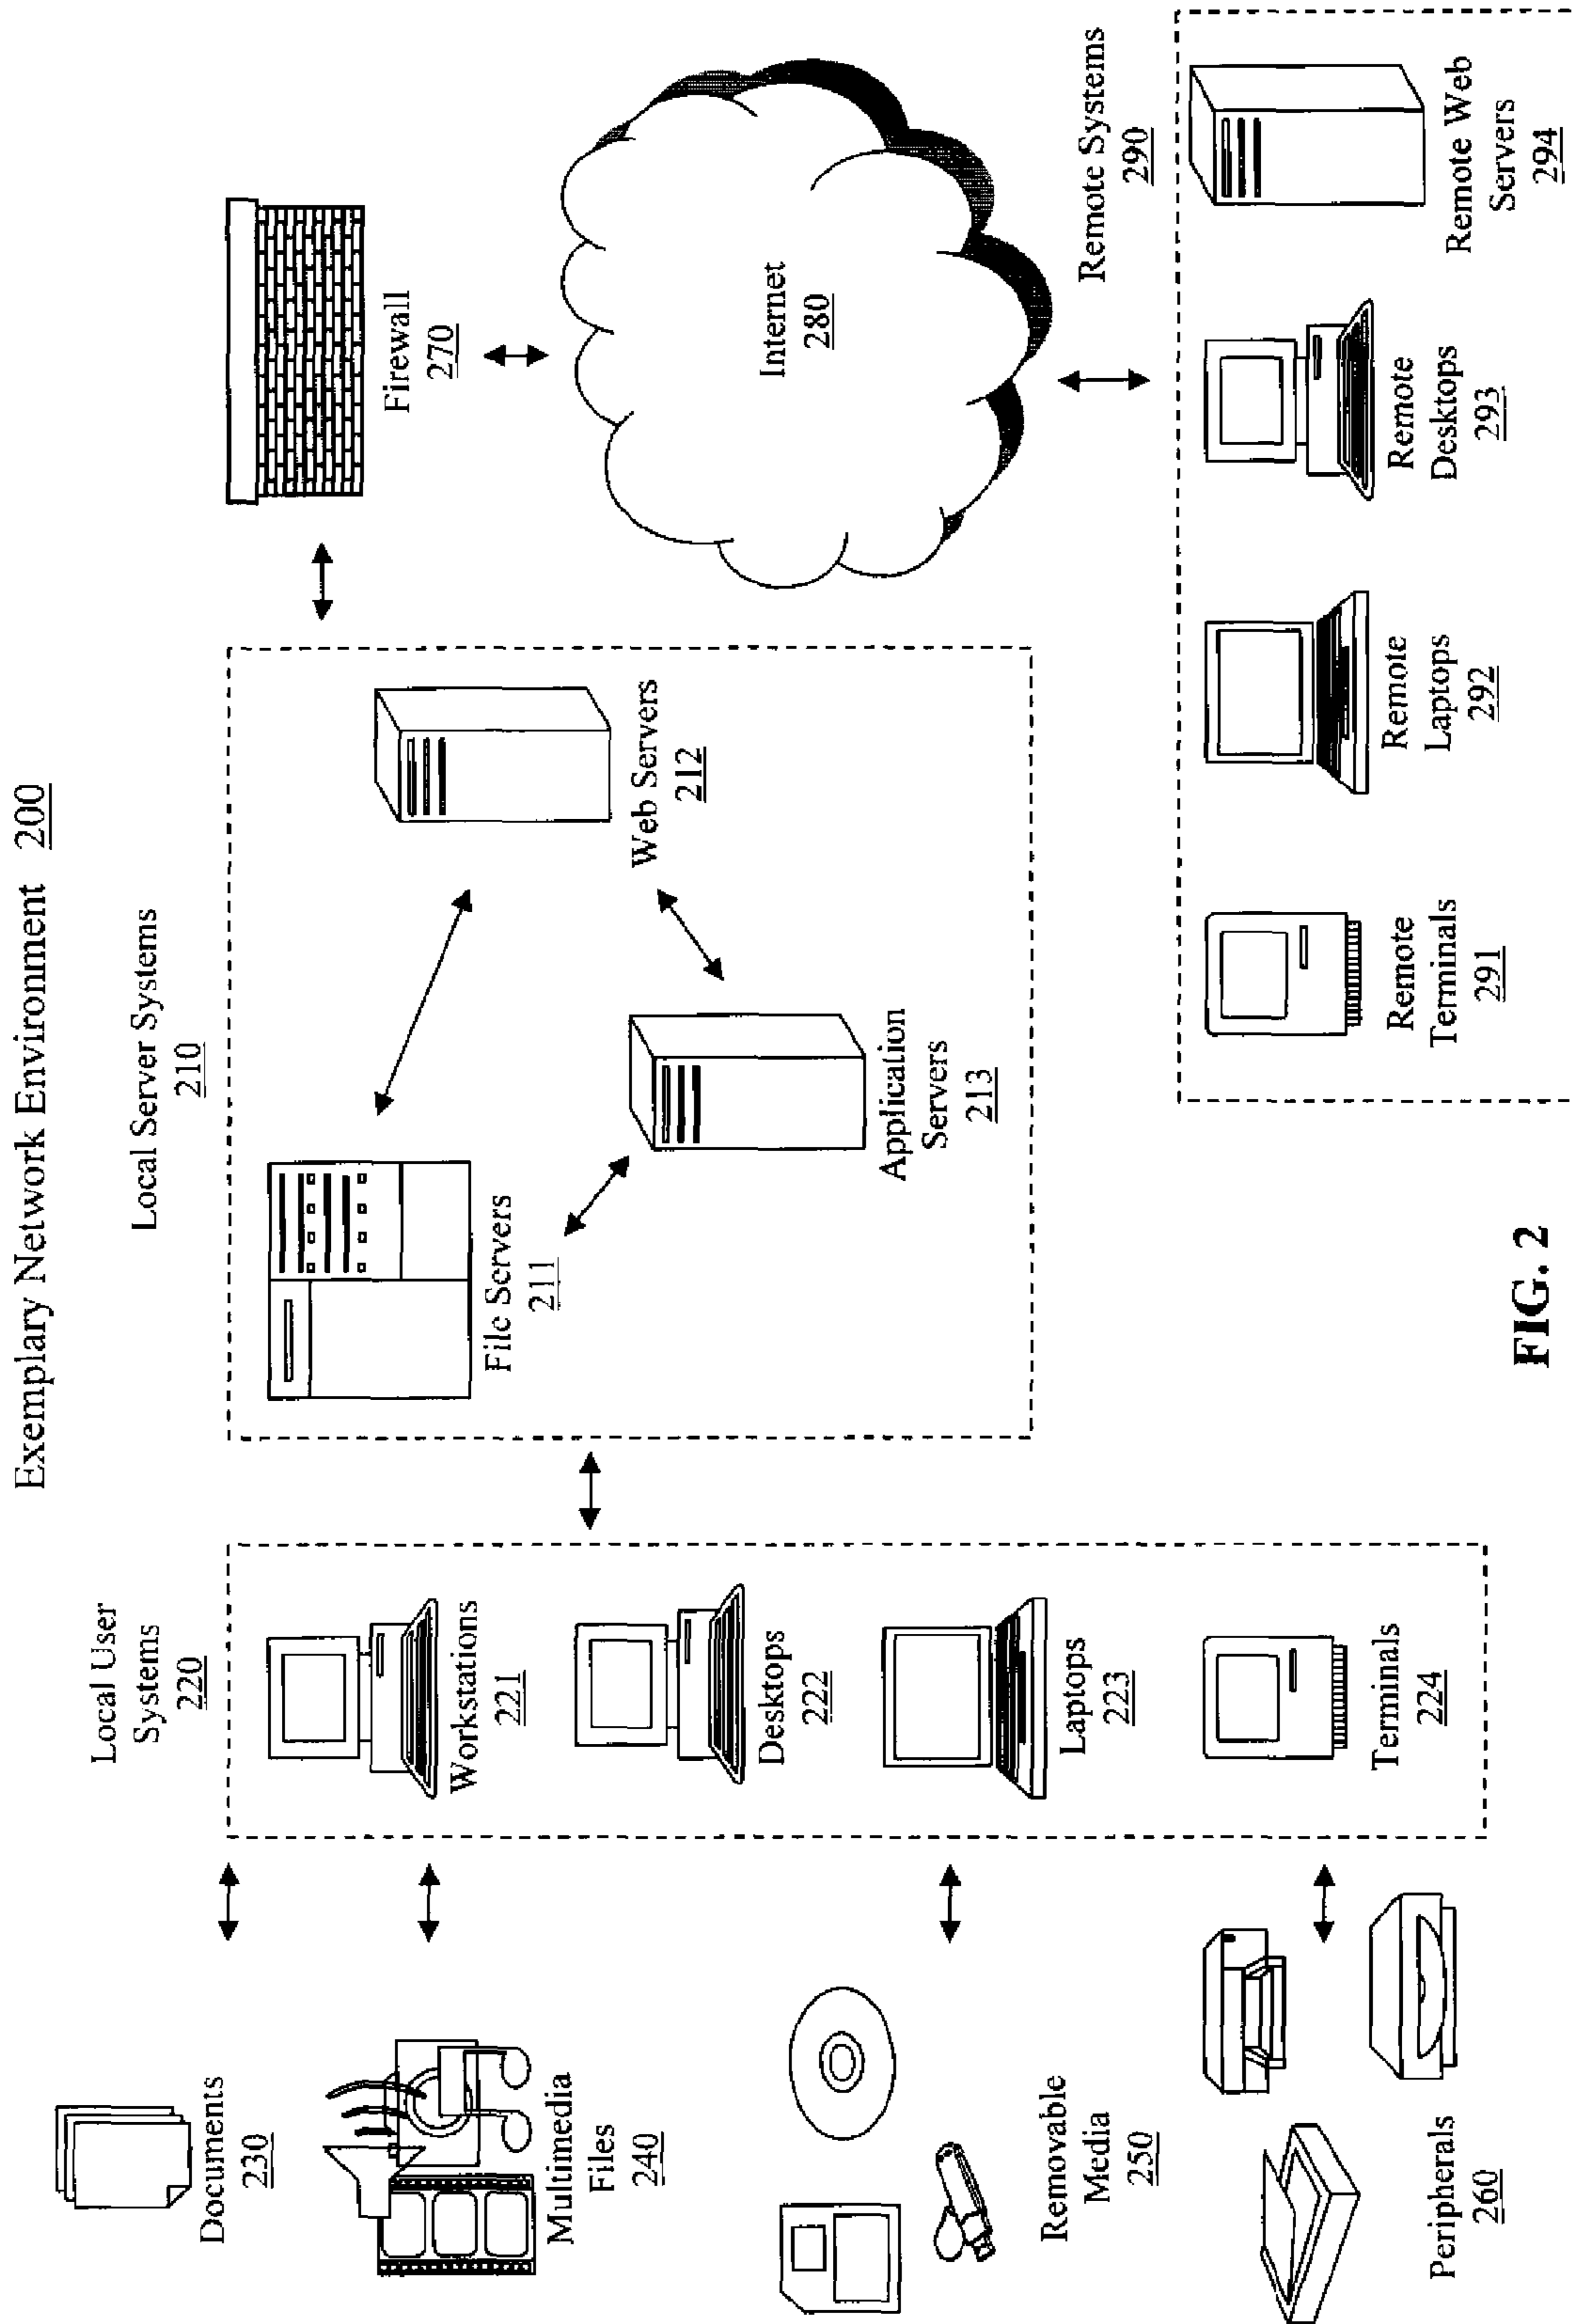


FIG. 2

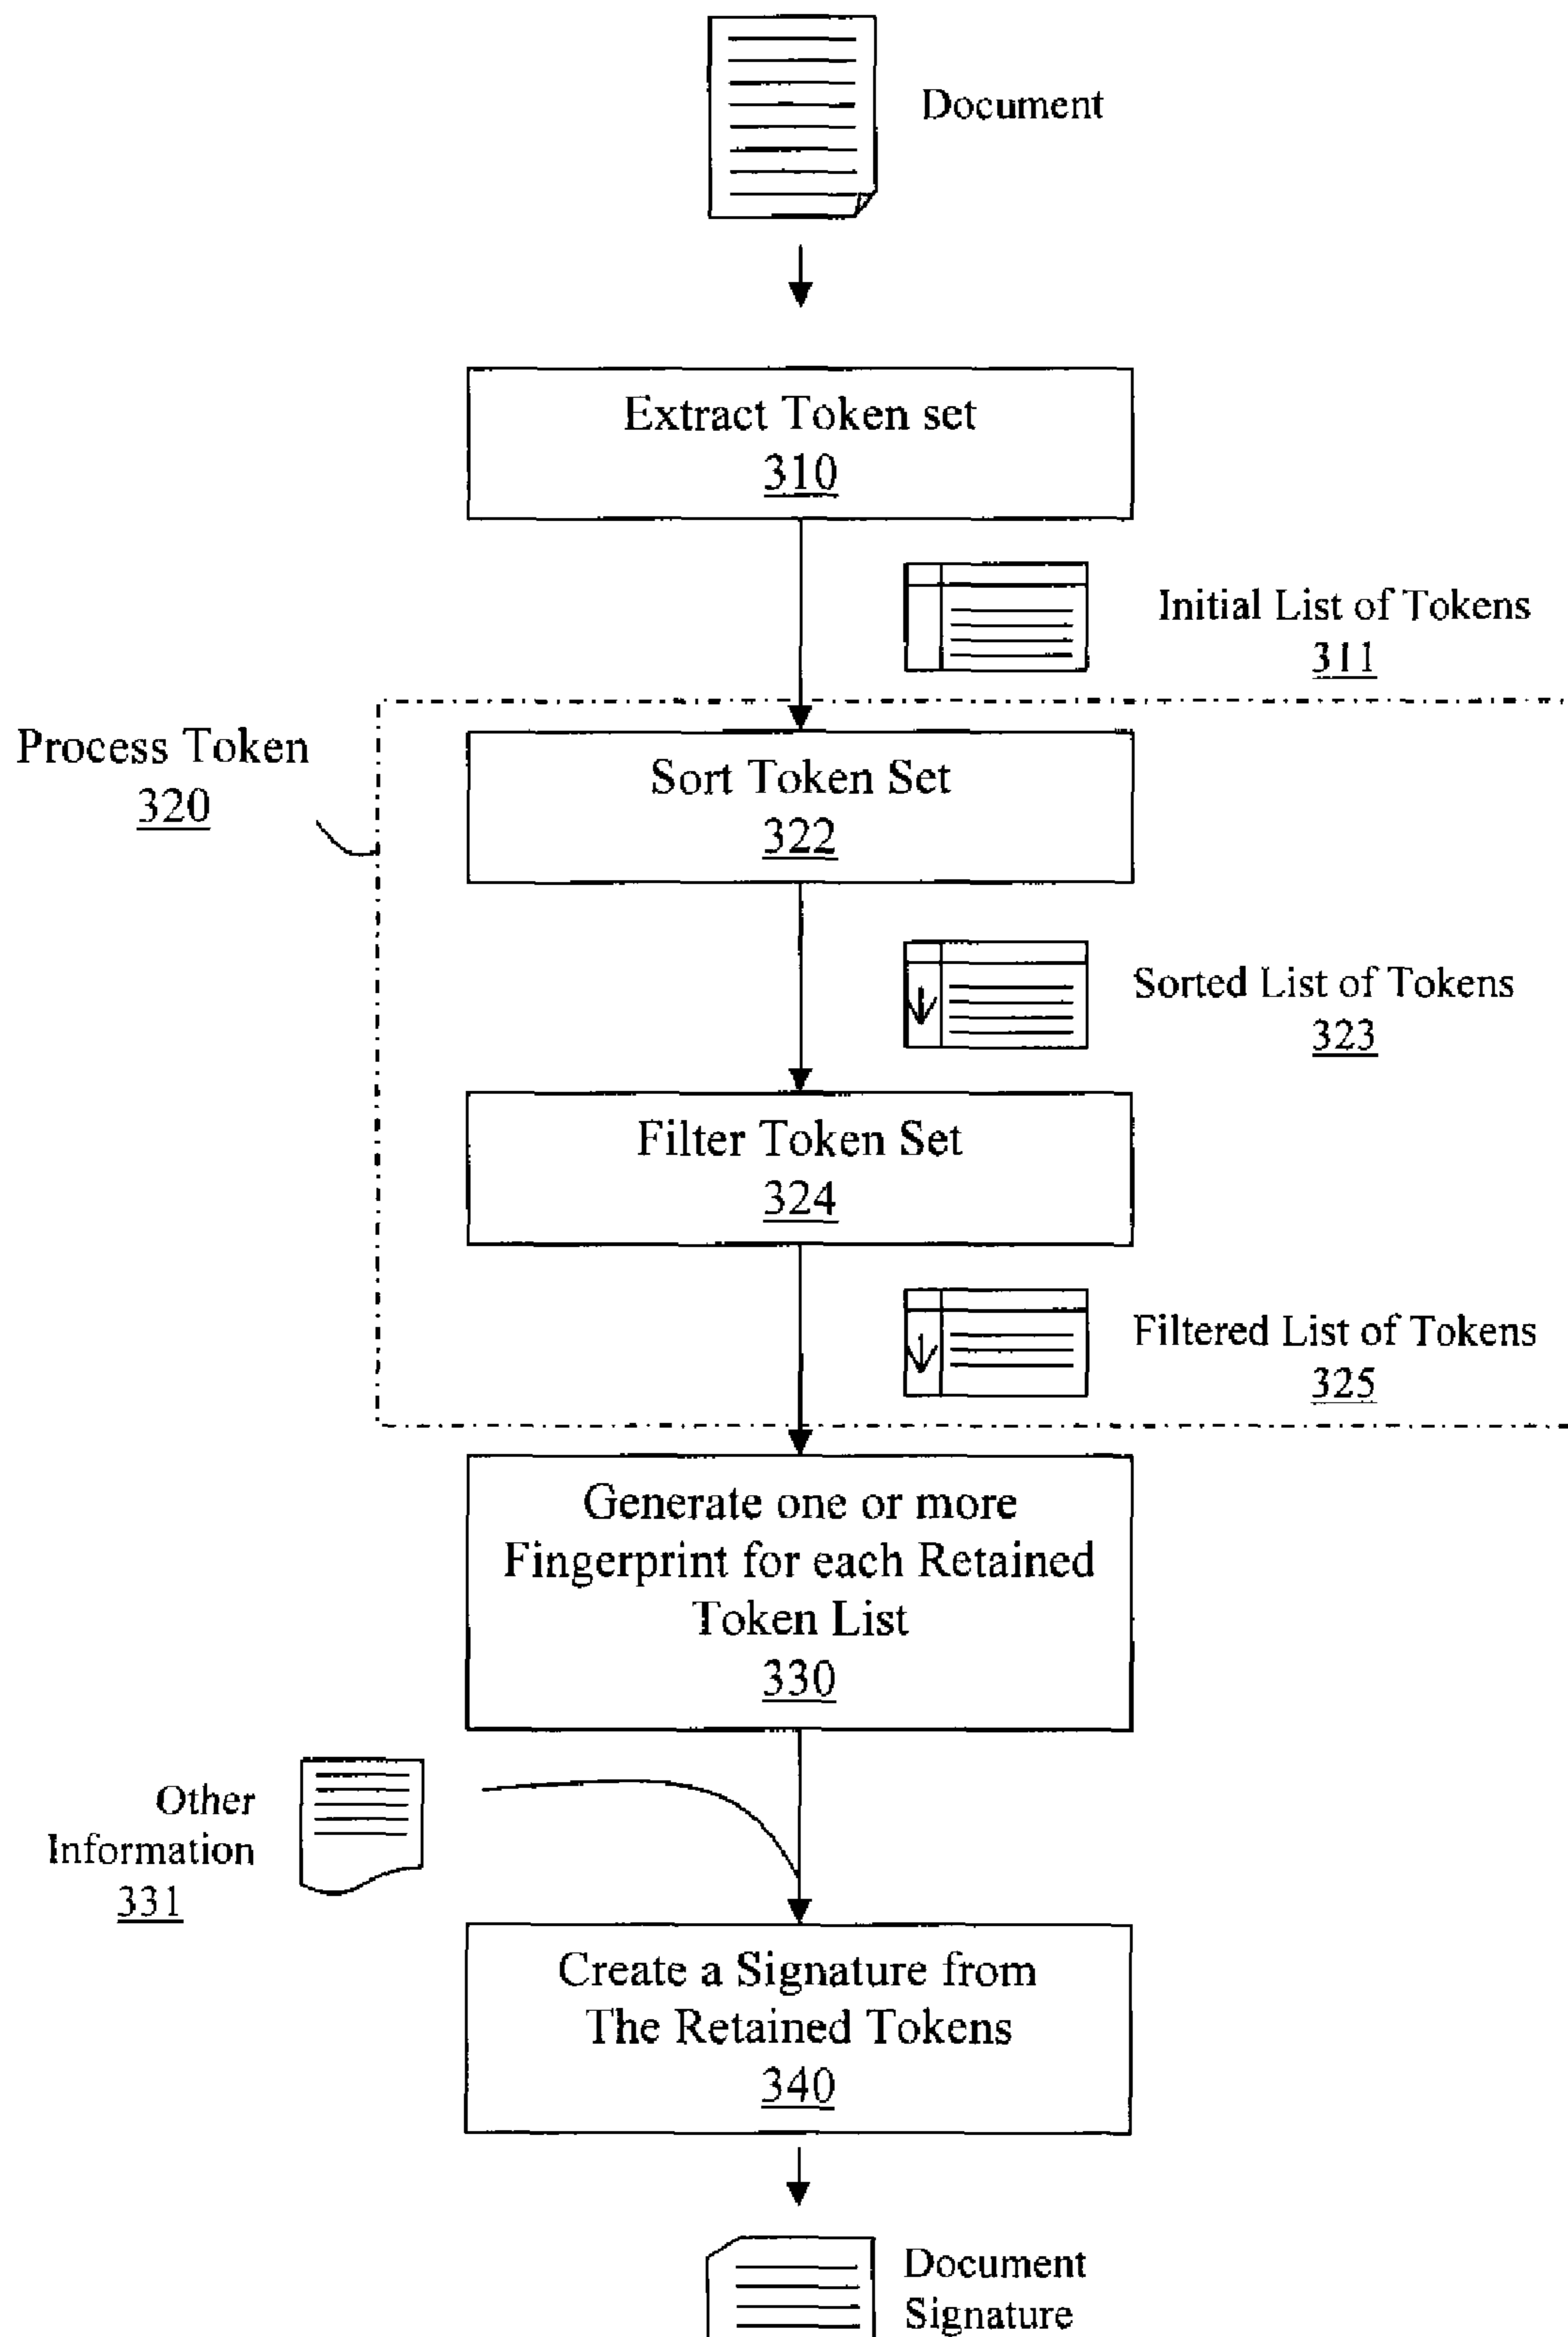


FIG. 3

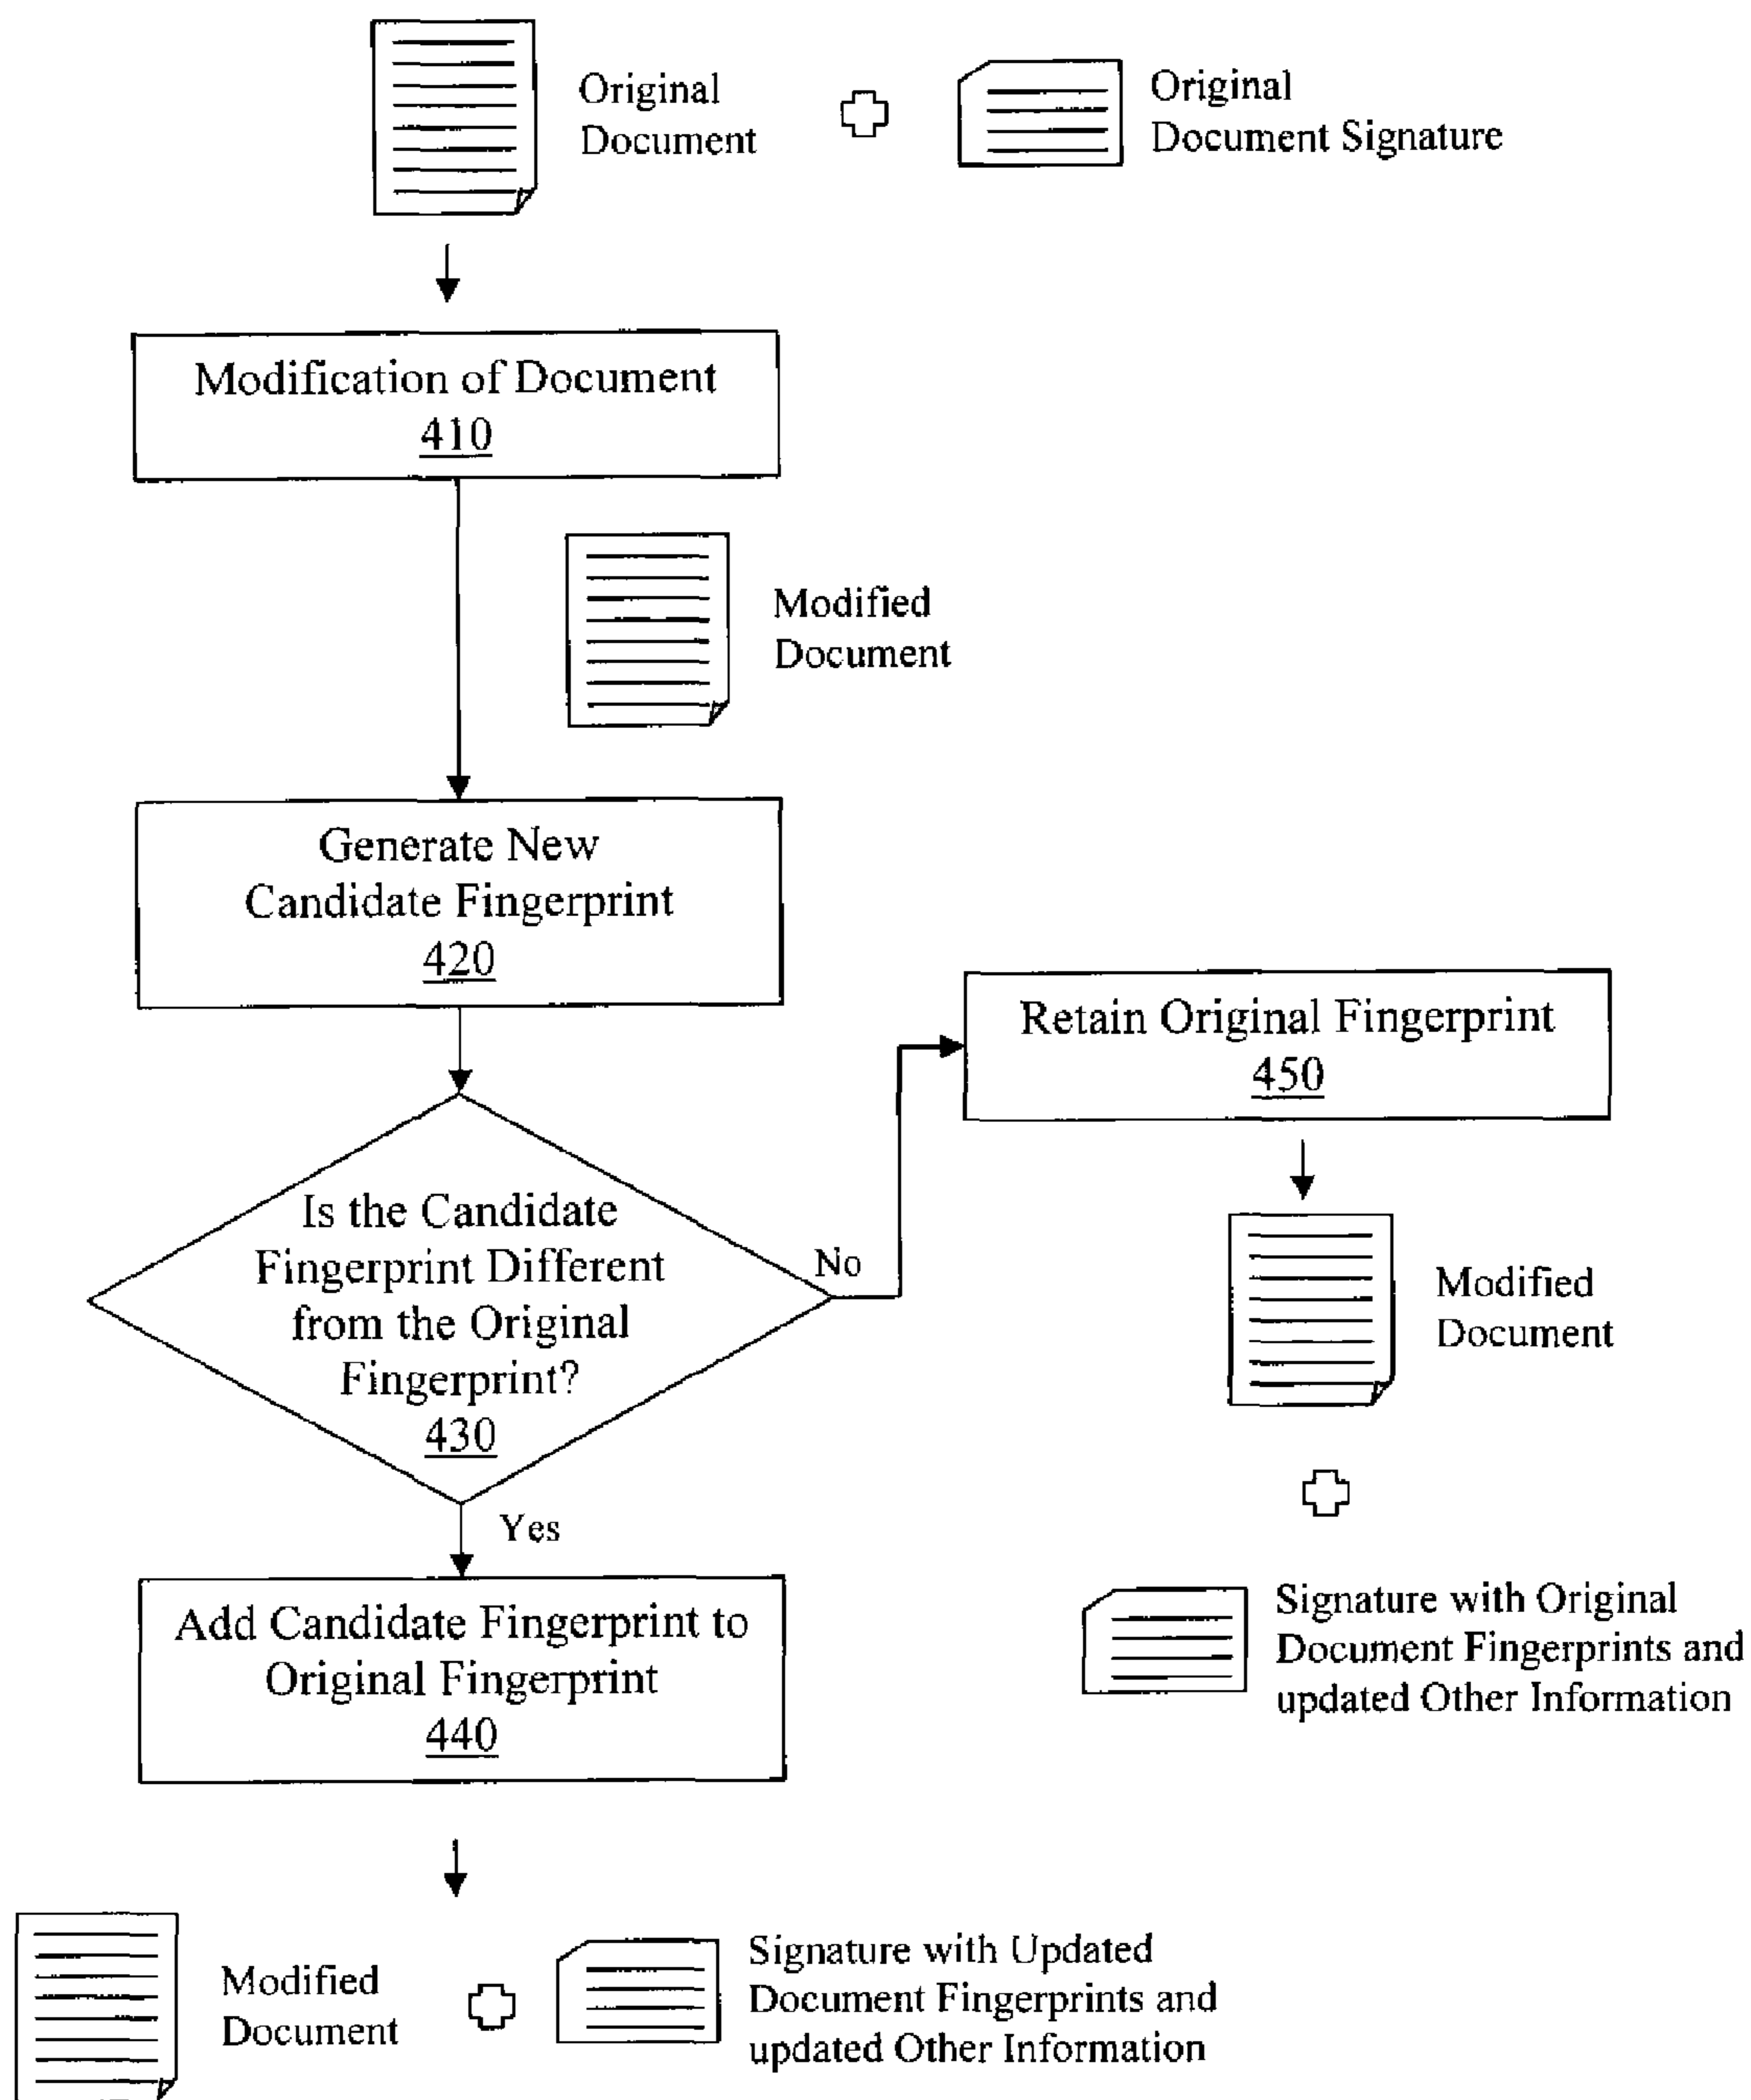


FIG. 4




FIG. 5




FIG. 6




FIG. 7

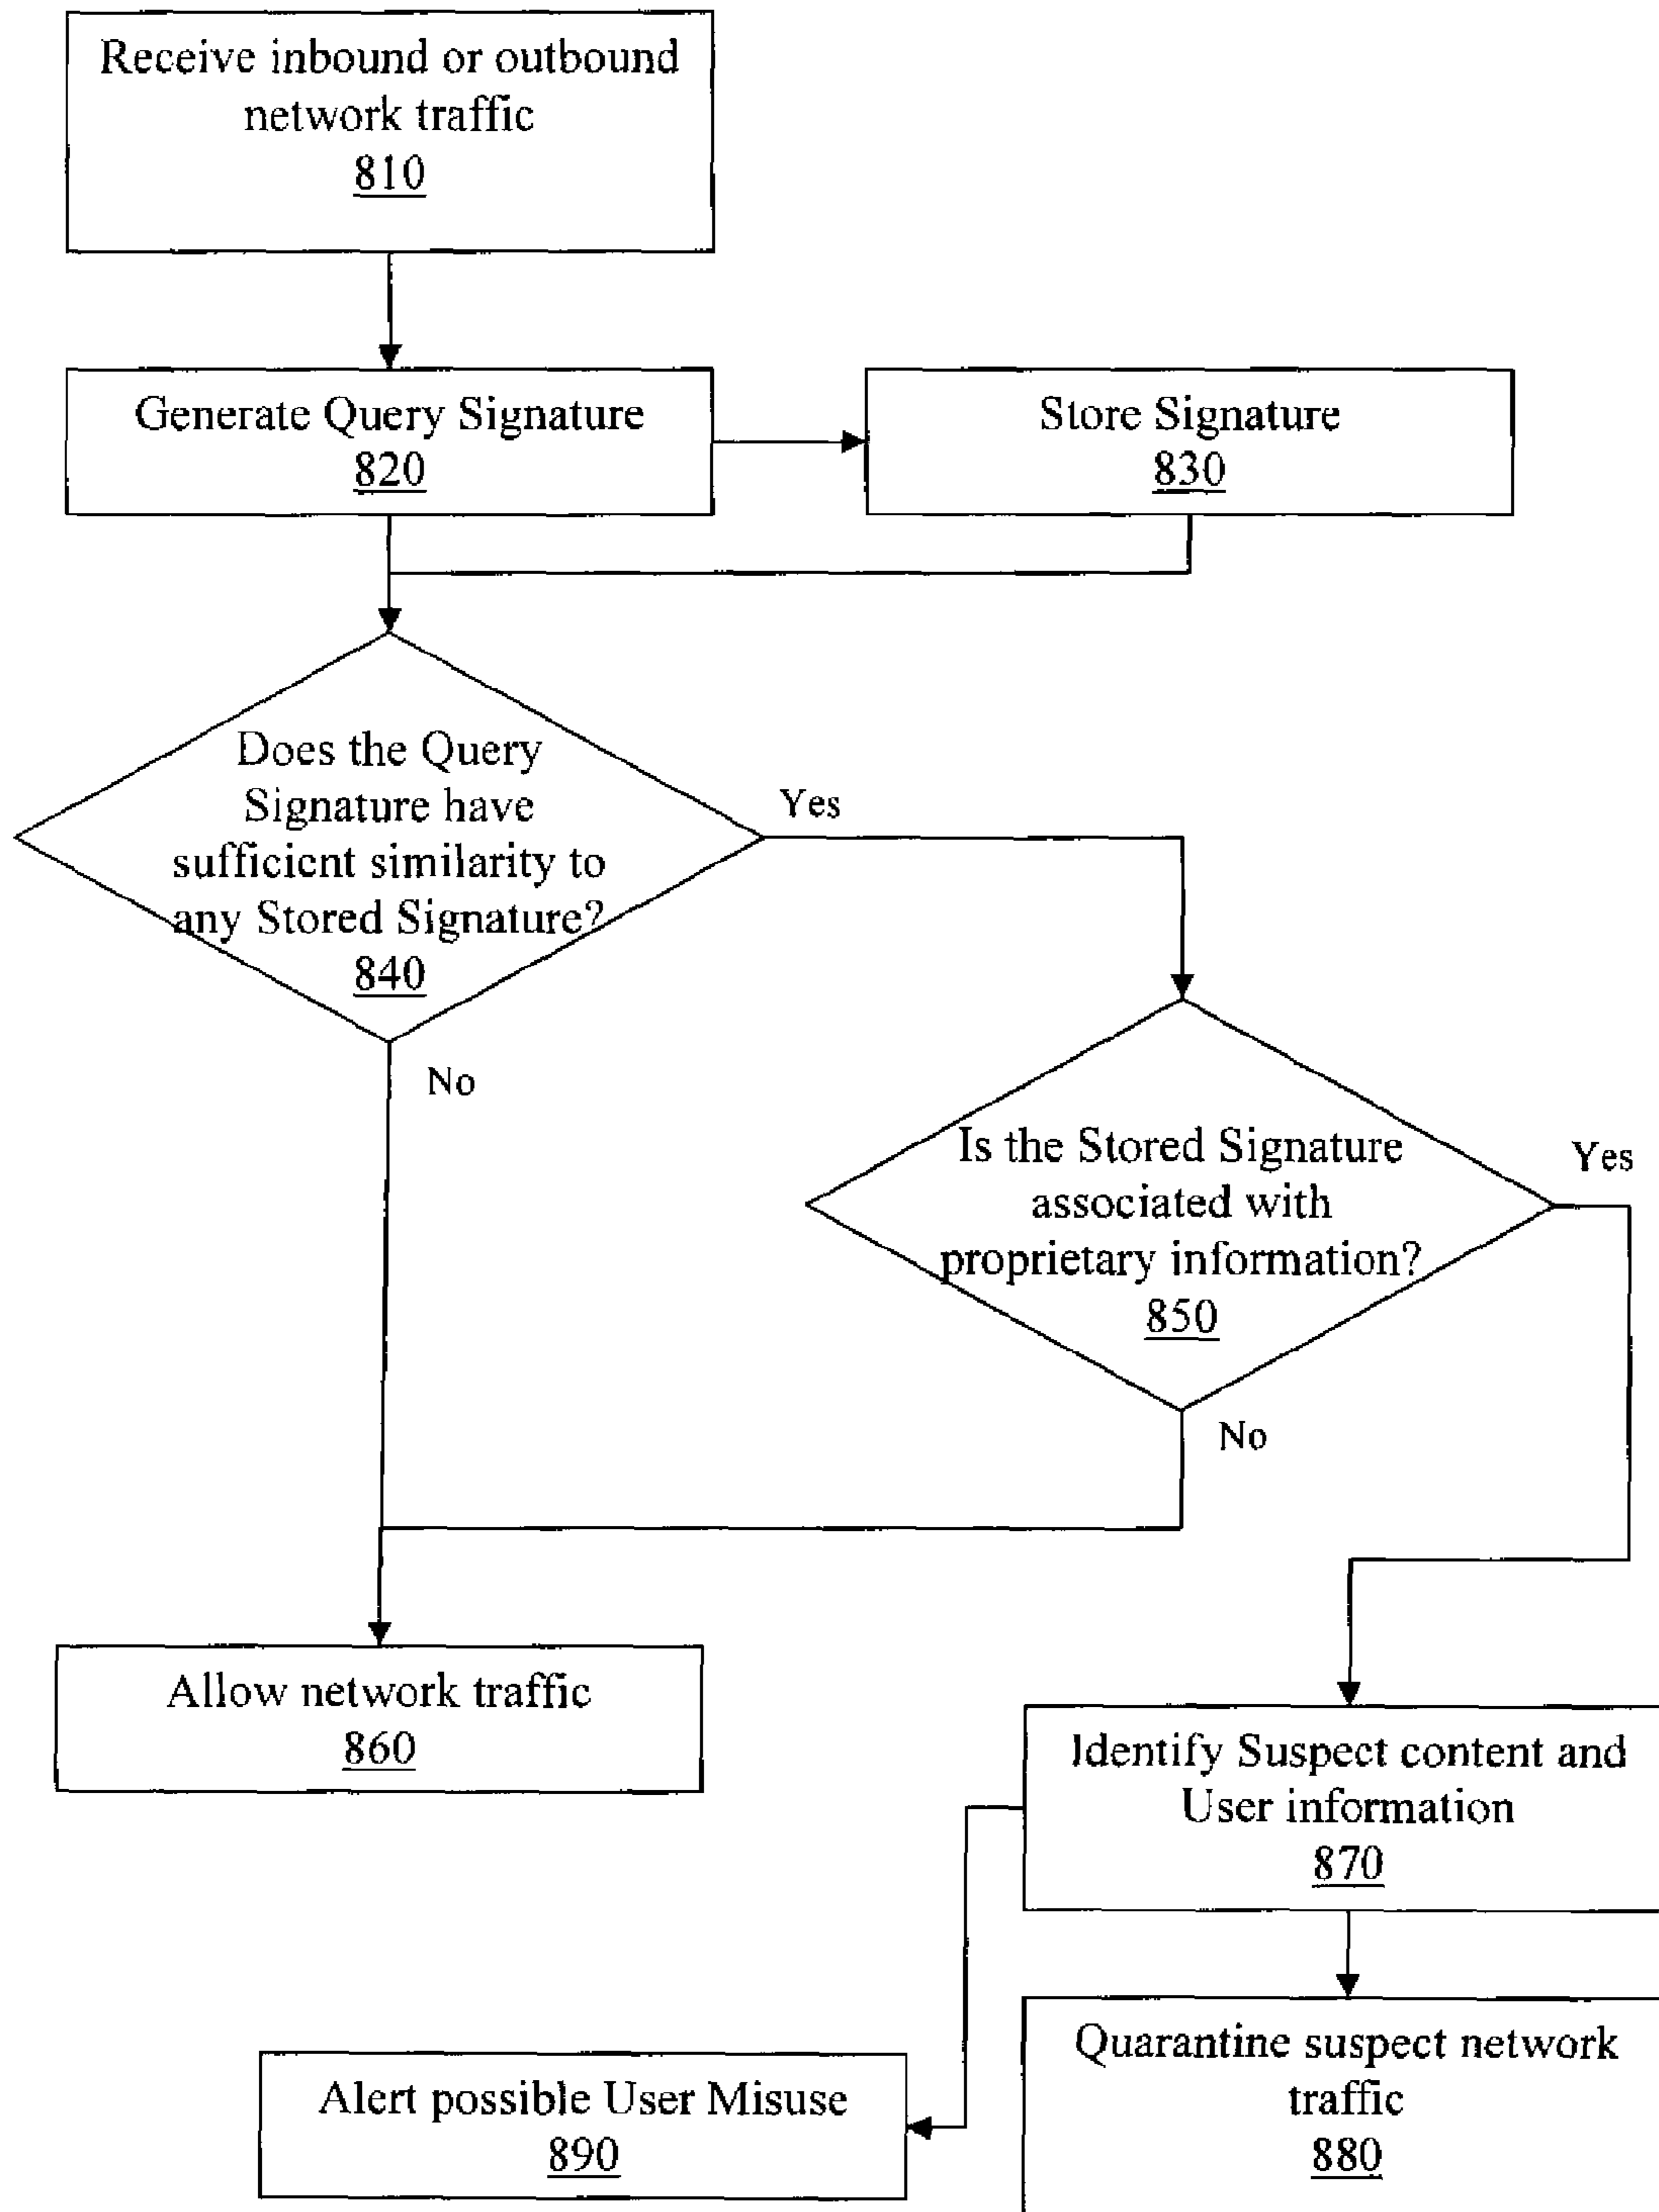


FIG. 8




FIG. 9

AUTOMATED FORENSIC DOCUMENT SIGNATURES

FIELD OF THE INVENTION

This invention relates generally to methods and systems for computer data management and tracking. Specifically, it relates to methods and systems of digital data identification and the creation, storage, management, processing and comparison of content sensitive digital signatures.

BACKGROUND OF THE INVENTION

Over the last decade, the use of computers and the Internet has grown exponentially. Indeed, for many individuals, government agencies and private corporations it is an integral part of their daily lives and business practices. People can communicate, transfer information, engage in commerce and expand their educational opportunities with little more than a few key strokes and the click of a mouse. Like revolutionary technologies before it, the great advancement of computer systems, information technology and the Internet carries enormous potential both for advancement and for abuse. Unfortunately, criminals exploit these same technologies to commit crimes and harm the safety, security, and privacy of the society.

Although there are no exact figures on the cost of computer crimes in America, estimates run into the billions of dollars each year. The United States Federal Bureau of Investigation (FBI) has indicated that digital evidence has spread from a few types of investigations, such as hacking and child pornography, to virtually every investigative classification, including fraud, extortion, homicide, identity theft, and so on. Although there are as yet no definitive statistics on the scope of the problem, there is no doubt that the number of crimes involving computers and the Internet is rising dramatically. A survey conducted by the Computer Security Institute in 2007 revealed substantial increases in computer crime. About half (46%) of the companies and government agencies surveyed reported a security incident within the preceding twelve months. The reported total loss of the participants is \$66,930,950. The average annual loss for each participant is \$350,424 compared to \$168,000 for the previous year. And unlike more traditional crimes, computer crime is especially difficult to investigate. Other criminal and terrorist acts and preparations leading to such acts, increasingly involve the use of computer systems and information technologies as well. These criminal and terrorist activities leave behind a trail of digital evidence. Digital evidence varies widely in formats and can include computer files, digital images, sound and videos, e-mail, instant messages, phone records, and so on. They are routinely gathered from seized hard drives, file servers, Internet data, mobile digital devices, digital cameras and numerous other digital sources that are growing steadily in sophistication and capacity.

Computer forensics is the practice of acquiring, preserving, analyzing, and reporting on data collected from a computer system, which can include personal computers, server computers, and portable electronic devices such as cellular phones, PDAs and other storage devices. Collecting and analyzing these types of data is usually called digital data identification. The goal of the process is to find evidence that supports or refutes some hypothesis regarding user activity on the system. When accurately and timely identified by a forensic investigator, digital evidence can provide the invaluable proof that helps the conviction of a criminal, or prevents a looming terrorist attack. A delay in identifying suspect data

occasionally results in the dismissal of some criminal cases, where the evidence is not being produced in time for prosecution.

The amount of digital evidence is growing rapidly. Not only has the number of crimes involving digital evidence increased dramatically over time, but the total volume of data that is involved has increased at an even faster pace. This is the result of the increased presence of digital devices at crime scenes combined with a heightened awareness of digital evidence by investigators. Given the declining prices of digital storage media and the corresponding increases in sales of storage devices, the volume of digital information that investigators must deal with is likely to continue its meteoric increase.

A typical computer forensic process involves first the determination that the evidence requirements merit a forensic examination. Individuals who are expected to have access to that evidence are then identified. Further, all computer systems used by these individuals which might contain relevant data are located. Forensic images of those systems are taken, and analyzed for relevant evidence. Traditionally, a forensic investigator seizes all storage media, creates a drive image or duplicates it, and then conducts their examination of the data on the drive image or duplicate copy to preserve the original evidence. A "drive image" is an exact replica of the contents of a storage device, such as a hard disk, stored on a second storage device, such as a network server or another hard disk. One of the first steps in the examination process is to recover latent data such as deleted files, hidden data and fragments from unallocated file space. Digital forensic analysis tools used today are stand alone systems that are not coordinated with systems used by the forensic investigators and Information Technology (IT) staff. Current computer forensics analysis is largely a manual labor intensive process. It requires computer forensic investigators that have specialized training. The cost of the analysis is high. The rate for some computer forensic investigators can be more than \$250/hour. It usually requires a long analysis time taking from days to weeks. Because it is a manual process, there is potential for human error resulting in missed data and missed discovery. In addition, when facing a complex investigation that involves a large number of computer systems, it is difficult to determine what systems to analyze. This may have two undesirable results: expending limited time and resources on useless systems, or missing systems that contain vital information.

The tremendous increase in data exacerbates these problems for forensic investigators. The number of pieces of digital media and their increasing size will push budgets, processing capability and physical storage space available to the forensic investigators to their limits. In an effort to reduce the volume of digital files for review, seized digital evidence is processed to reduce the amount of this data. Presently, there is no effective means to quickly sort through the amount of data based on the content of the data, and identify documents and files of interest for further detailed examination. Present solutions still require manual review from forensic investigators to identify specific data needed to prove guilt or innocence.

Government and business entities use sophisticated computers systems to store, track and disseminate information within the entity and communicate with outside individuals and entities. Information can be stored as files that exist on a computer file system, and can exist in many heterogeneous forms such as plain text documents, formatted documents (e.g. Microsoft Word® documents, Open Document Format documents), spread sheets, presentations, Portable Document Format documents, images of paper documents, graphics, sound recordings, videos, faxes, email messages, voice mes-

sages, web pages, and other stored digital media. Information can also be stored as entries in databases such as a relational database or a document management system. This information is subject to a wide range of user manipulations, such as create, edit, copy, rename, move, delete and backup. Information can also move among the entity computer systems through various communication means, such as emails, attachments, file sharing, shared file systems and push technology. Information can also leave the entity computer systems either by someone within the entity sending it to an outsider, or can be retrieved by an outsider from the entity computer systems by obtaining information containing removable storage media or through network access protocols such as HTTP, FTP, and peer-to-peer file sharing. All of this creation, manipulation, transfers, and communication of digital information can be part of the legitimate business process. However, abuse of the computer system also involves the same processes of creation, manipulation, transfer, and communication of information, albeit unauthorized or illegitimately. The Computer Security Institute 2007 survey also revealed that insider abuse of the network access or email edged out virus incidents as the most prevalent security problem. While a majority of all computer attacks enter via the Internet, the most significant of all dollar losses stem from internal intrusions.

The most important asset of many companies is their Intellectual Property (IP). Customer lists, customer credit card lists, copyrights including computer code, confidential product designs, proprietary information such as new products in development, and trade secrets are all forms of IP that can be used against the company by its competitors. Common risks for a corporation may be theft of trade secrets and other privileged information, theft of customer or partner information, disclosure of confidential information, and disclosure of trade secrets and other valuable information (designs, formulas etc.).

Corporations may also incur liability or exposure to risks when unauthorized contents are stored in the computer systems, such as child pornographic material, or pirated copies of media or software. An organization must know which of its assets require protection and the real and perceived threats against them.

Current information security builds layers of firewalls and content security at the network perimeter, and utilizes permissions and identity management to control access by trusted insiders to digital assets, such as business transactions, data warehouses and files. This structure lulls the business managers into a false sense of security. Many employees are restricted in their access to sensitive data, but access control is usually not easily fine tuned to accommodate the ever changing assignments and business needs of all the employees. Moreover, as is necessary to perform their function, Information Technology (IT) employees have access to sensitive data and processes. Indeed, IT employees are the custodians and authors of those objects. This may place them in positions to reveal information to others that will damage the company or directly sabotage a company's operations in various ways. IT employees who are disgruntled, angry, or seeking to steal information for profitable gain, may attempt to steal sensitive digital information which could lead to substantial losses for the organization. A laid-off employee is a prime source of potential leakage of such information.

Content-security tools based on HTTP/SMTP proxies are used against viruses and spam. However, these tools weren't designed for intrusion prevention. They don't inspect internal traffic; they scan only authorized e-mail channels. They rely on file-specific content recognition and have scalability and

maintenance issues. When content security tools don't fit, they are ineffective. Relying on permissions and identity management is like running a retail store that screens you coming in but doesn't put magnetic tags on the clothes to prevent you from wearing that expensive hat going out.

A hash analysis is a method that can be used for comparing the content of digital evidence. A cryptographic one-way hash (or "hash" for short) can be a way to calculate a digital fingerprint: a very large number that often uniquely identifies a digital file. A hash is a calculated function on the bits that make up a file. Therefore, two files with different names but the exact same contents will produce the same hash. However, using hash systems to identify conclusive or known suspect files faces several challenges. By design of the hash function, a small difference, even a single bit, in the input file will generate a significantly different output hash. The difference between two hash numbers does not reflect the level of similarity of the input files. The hash method cannot be used to identify files that have been altered, whether minimally or substantially. They are therefore not able to identify derivative files, files that contain common contents but are arranged or formatted differently or contain more or less other content. For the same reason, hash analysis is not effective against multimedia files (image, video, and sound). As a consequence, an individual using these files to commit crimes may escape hash based detection and prosecution.

It would be beneficial and desirable to integrate newer, advanced technologies to automate the detection and classification process for suspect files and identify related altered or derivative files. This would allow forensic investigators to focus on identifying relevant data during the forensic process and addresses many of the problems of efficiency, cost and delay facing digital forensic examinations today. There is also a need for a technology to scan and manage digital data on a computer system based on the content of the data. There is a further need for a solution to allow government agencies and corporations to automatically monitor and prevent unauthorized use or exchange of classified or proprietary data.

SUMMARY OF THE INVENTION

The present invention is a method, system, and computer readable media for proactively generating, preserving and comparing computer forensic evidence for a computer system. The method involves generating at least one signature for at least one target based on the content of the target. The at least one signature can be generated at any time, or when a predetermined operation is commenced. The at least one generated signature can be stored, or not, prior to or after forensic use. The generated signature(s) are compared with one or more previously generated signature(s) to determine whether any compared signatures have similarities above a predetermined threshold. Alternatively, the present invention could, at any time, simply compare previously existing signatures generated from a target.

The target can be any file, any file that is owned by a user, any operating system file, any file that is part of a proprietary information system, or any file that is related to a network intrusion attack. When the target is any type of file, the predetermined operation can be any one or more of creating, deleting, renaming, editing, moving, updating, linking, merging, modifying and copying the file. The target could also be a database entry; and when a database entry, the predetermined operation can be any one or more of selecting, inserting, updating, deleting, merging, beginning work, committing, rollback, creating, dropping, truncating, and altering of the database entry. The target can further be a database defi-

5

When the target is a database definition, the predetermined operation can be any one or more of creating, dropping and altering the database definition.

The target can also be network traffic; and when network traffic, the predetermined operation can be the occurrence of network traffic entering a network or leaving a network, or a network traffic is initiated from a computer system, or a computer system receives network traffic. The network traffic may be any one or more of a signal protocol, an email, an attachment of an email, an instant message conversation, a text message, a remote login, a virtual private network, a viewed webpage, a file transfer and file sharing.

Generating the at least one signature can involve extracting a set of tokens from the at least one target, processing the set of tokens, generating a fingerprint from the set of tokens, and generating the signature for the target by combining the fingerprints with other related information of the target. Processing the set of tokens can include sorting the set of tokens, and may further include filtering the set of tokens. The method for generating the fingerprints may involve a hash method, or an implementation of a bit vector method.

Other related information of the target can be accessible by an operating system, and can be any one or more of file name, date of record, time of record, user or owner information, network address, network protocol, access history and fingerprint history. Other related information of the target could also be information accessible by an application.

The generated signature(s) could be stored in a manner preventing deletion or modification by a user, other than a user with special access rights, such authorized personnel or a forensic investigator. The signature(s) could further be made available only to authorized personnel or a forensic investigator with access rights. The signature(s) and respective targets can be stored on the same computer system, different computer systems, and/or on a shared file system. Finally, the signature(s) can be stored on write-once, read-many media.

In another aspect of the present invention, a computer readable medium is provided that configures a computer system to perform the methods described above of proactively generating, preserving and comparing computer forensic evidence for a computer system. In summary, computer readable medium facilitates the method of generating at least one signature for at least one target based on the content of the target; and comparing the at least one generated signature with at least one previously generated signature to determine whether the signatures have similarities above a predetermined threshold.

In a further aspect of the present invention, the present invention also provides an apparatus for the generation, preservation and comparison of computer forensic evidence. The apparatus/system can include a processor arranged to generate at least one signature for at least one target based on the content of the target, and a comparator configured to compare the at least one generated signature with at least one previously generated signature to determine whether the signatures have similarities above a predetermined threshold. In addition, the system can additionally include an extension module configured to trigger signature generation upon occurrence of a certain action, and a mechanism for storing the generated signatures. The implemented system may have an operating system service (e.g., a Windows® service or Unix/Linux daemon) running in the background to generate a signature for a given file and to store it, and then to query the stored signatures to determine similarity with other signatures.

6

DESCRIPTION OF THE DRAWINGS

FIG. 1 is a schematic diagram of an exemplary computing environment;

FIG. 2 is a schematic diagram of an exemplary network environment;

FIG. 3 is a flow chart illustrating an exemplary method for generating a signature for a document;

FIG. 4 is a flow chart illustrating document modification and new fingerprint generation pursuant to one embodiment of the present invention;

FIG. 5 is a flow chart illustrating an exemplary method to perform a latent signature

FIG. 6 is a flow chart illustrating an exemplary method for user misuse detection;

FIG. 7 is a flow chart illustrating another exemplary method for user misuse detection through the use of user signature profiles;

FIG. 8 is a flow chart illustrating an exemplary method for the detection of an unauthorized network communication of sensitive information; and

FIG. 9 is a schematic block diagram illustrating an exemplary embodiment of a system of the present invention, showing event trigger, fingerprint/signature generation, signature query and comparison, and signature storage.

DETAILED DESCRIPTION

Example Computing Environment

FIG. 1 and the following discussion are intended to provide a brief general description of a suitable computing environment in which an example embodiment of the invention may be implemented. It should be understood, however, that handheld, portable, and other computing devices of all kinds are contemplated for use in connection with the present invention. While a general purpose computer is described below, this is but one example. The present invention also may be operable on a thin client having network server interoperability and interaction. Thus, an example embodiment of the invention may be implemented in an environment of networked hosted services in which very little or minimal client resources are implicated, e.g., a networked environment in which the client device serves merely as a browser or interface to the World Wide Web.

Although not required, the invention can be implemented via an application programming interface (API), for use by a developer or tester, and/or included within the network browsing software which will be described in the general context of computer-executable instructions, such as program modules, being executed by one or more computers (e.g., client workstations, servers, or other devices). Generally, program modules include routines, programs, objects, components, data structures and the like that perform particular tasks or implement particular abstract data types. Typically, the functionality of the program modules may be combined or distributed as desired in various embodiments. Moreover, those skilled in the art will appreciate that the invention may be practiced with other computer system configurations. Other well known computing systems, environments, and/or configurations that may be suitable for use with the invention include, but are not limited to, personal computers (PCs), server computers, hand-held or laptop devices, multi-processor systems, microprocessor-based systems, programmable consumer electronics, network PCs, minicomputers, mainframe computers, and the like. An embodiment of the invention may also be practiced in distributed computing environments where tasks are performed by remote processing

devices that are linked through a communications network or other data transmission medium. In a distributed computing environment, program modules may be located in both local and remote computer storage media including memory storage devices.

FIG. 1 thus illustrates an example of a suitable computing system environment 100 in which the invention may be implemented, although as made clear above, the computing system environment 100 is only one example of a suitable computing environment and is not intended to suggest any limitation as to the scope of use or functionality of the invention. Neither should the computing environment 100 be interpreted as having any dependency or requirement relating to any one or a combination of components illustrated in the exemplary operating environment 100.

With reference to FIG. 1, an example system for implementing the invention includes a general purpose computing device in the form of a computer 110. Components of the computer 110 may include, but are not limited to, a processing unit 120, a system memory 130, and a system bus 121 that couples various system components including the system memory to the processing unit 120. The system bus 121 may be any of several types of bus structures including a memory bus or memory controller, a peripheral bus, and a local bus using any of a variety of bus architectures. By way of example, and not limitation, such architectures include Industry Standard Architecture (ISA) bus, Micro Channel Architecture (MCA) bus, Enhanced ISA (EISA) bus, Video Electronics Standards Association (VESA) local bus, Peripheral Component Interconnect (PCI) bus (also known as Mezzanine bus), and PCI-Express bus.

The computer 110 typically includes a variety of computer readable media. Computer readable media can be any available media that can be accessed by the computer 110 and includes both volatile and nonvolatile, removable and non-removable media. By way of example, and not limitation, computer readable media may comprise computer storage media and communication media. Computer storage media includes both volatile and nonvolatile, removable and non-removable media implemented in any method or technology for storage of information such as computer readable instructions, data structures, program modules or other data. Computer storage media includes, but is not limited to, random access memory (RAM), read-only memory (ROM), Electrically-Erasable Programmable Read-Only Memory (EEPROM), flash memory or other memory technology, compact disc read-only memory (CDROM), digital versatile disks (DVD) or other optical disk storage, magnetic cassettes, magnetic tape, magnetic disk storage or other magnetic storage devices, or any other medium which can be used to store the desired information and which can be accessed by the computer 110. Communication media typically embodies computer readable instructions, data structures, program modules or other data in a modulated data signal such as a carrier wave or other transport mechanism and includes any information delivery media. The term "modulated data signal" means a signal that has one or more of its characteristics set or changed in such a manner as to encode information in the signal. By way of example, and not limitation, communication media includes wired media such as a wired network or direct-wired connection, and wireless media such as acoustic, radio frequency (RF), infrared, and other wireless media. Combinations of any of the above should also be included within the scope of computer readable media.

The system memory 130 includes computer storage media in the form of volatile and/or nonvolatile memory such as ROM 131 and RAM 132. A basic input/output system 133

(BIOS), containing the basic routines that help to transfer information between elements within computer 110, such as during start-up, is typically stored in ROM 131. RAM 132 typically contains data and/or program modules that are immediately accessible to and/or presently being operated on by the processing unit 120. By way of example, and not limitation, FIG. 1 illustrates operating system 134, application programs 135, other program modules 136, and program data 137. RAM 132 may contain other data and/or program modules.

The computer 110 may also include other removable/non-removable, volatile/nonvolatile computer storage media. By way of example only, FIG. 1 illustrates a hard disk drive 141 that reads from or writes to non-removable, nonvolatile magnetic media, a magnetic disk drive 151 that reads from or writes to a removable, nonvolatile magnetic disk 152, and an optical disk drive 155 that reads from or writes to a removable, nonvolatile optical disk 156, such as a CD ROM or other optical media. Other removable/non-removable, volatile/nonvolatile computer storage media that can be used in the example operating environment include, but are not limited to, magnetic tape cassettes, flash memory cards, digital versatile disks, digital video tape, solid state RAM, solid state ROM, and the like. The hard disk drive 141 is typically connected to the system bus 121 through a non-removable memory interface such as interface 140, and magnetic disk drive 151 and optical disk drive 155 are typically connected to the system bus 121 by a removable memory interface, such as interface 150.

The drives and their associated computer storage media discussed above and illustrated in FIG. 1 provide storage of computer readable instructions, data structures, program modules and other data for the computer 110. In FIG. 1, for example, the hard disk drive 141 is illustrated as storing operating system 144, application programs 145, other program modules 146, and program data 147. Note that these components can either be the same as or different from operating system 134, application programs 135, other program modules 136, and program data 137. Operating system 144, application programs 145, other program modules 146, and program data 147 are given different numbers here to illustrate that, at a minimum, they are different copies. A user may enter commands and information into the computer 110 through input devices such as a keyboard 162 and pointing device 161, commonly referred to as a mouse, trackball or touch pad. Other input devices (not shown) may include a microphone, joystick, game pad, satellite dish, scanner, or the like. These and other input devices are often connected to the processing unit 120 through a user input interface 160 that is coupled to the system bus 121, but may be connected by other interface and bus structures, such as a parallel port, game port or a universal serial bus (USB).

A monitor 191 or other type of display device is also connected to the system bus 121 via an interface, such as a video interface 190. In addition to monitor 191, computers may also include other peripheral output devices such as speakers and a printer (not shown), which may be connected through an output peripheral interface 195.

The computer 110 may operate in a networked environment using logical connections to one or more remote computers, such as a remote computer 180. The remote computer 180 may be a personal computer, a server, a router, a network PC, a peer device or other common network node, and typically includes many or all of the elements described above relative to the computer 110, although only a memory storage device 181 has been illustrated in FIG. 1. The logical connections depicted in FIG. 1 include a local area network (LAN)

171 and a wide area network (WAN) 173, but may also include other networks. Such networking environments are commonplace in offices, enterprise-wide computer networks, intranets and the Internet.

When used in a LAN networking environment, the computer 110 is connected to the LAN 171 through a network interface or adapter 170. When used in a WAN networking environment, the computer 110 typically includes means for establishing communications over the WAN 173, such as the Internet. In a networked environment, program modules depicted relative to the computer 110, or portions thereof, may be stored in the remote memory storage device. By way of example, and not limitation, FIG. 1 illustrates remote application programs 185 as residing on a memory device 181. Remote application programs 185 include, but are not limited to web server applications such as Microsoft® Internet Information Services (IIS)® and Apache HTTP Server which provides content which resides on the remote storage device 181 or other accessible storage device to the World Wide Web. It will be appreciated that the network connections shown are exemplary and other means of establishing a communications link between the computers may be used.

One of ordinary skill in the art can appreciate that a computer 110 or other client devices can be deployed as part of a computer network. In this regard, the present invention pertains to any computer system having any number of memory or storage units, and any number of applications and processes occurring across any number of storage units or volumes. An embodiment of the present invention may apply to an environment with server computers and client computers deployed in a network environment, having remote or local storage. The present invention may also apply to a standalone computing device, having programming language functionality, interpretation and execution capabilities.

Example Network Environment

FIG. 2 illustrates an embodiment of a network environment in which an embodiment of the present invention can be implemented. The network environment 200 contains a number of local server systems 210, which may include a number of file servers 211, web servers 212, and application servers 213 that are owned and managed by the owner of the local network. These servers are in communication with local user systems 220 which may include a large variety of systems such as workstations 221, desktop computers 222, laptop computers 223, and thin clients or terminals 224. The local user systems 220 may contain their own persistent storage devices such as in the case of workstations 221, desktop computers 222, and laptop computers 223. They can also have access to the persistent storage provide by the local servers 210. In the case of thin clients and terminals 224, network storage may be the only available persistent storage. The local user systems are usually connected to a variety of peripherals 260 that handle data input and output, such as scanners, printers and optical drives. There may also be a number of different kinds of removable media 250 that attach to the user systems 220 at times. These removable media 250 can be based on magnetic recording, such as floppy disks and portable hard drives, or be based on optical recording, such as compact disks or digital video disks. Further, removable media can also be based on non-volatile memory such as flash memory which can be a USB flash drive, and all forms of flash memory cards. The users within the local network usually get access to the wider area network such as the Internet 280 though the local server systems 210 and typically some network security measures such as a firewall 270. There might also be a number of remote systems 290 that can be in communication with the local server systems 210 and also the

local user systems 220. The remote computer systems can be a variety of remote terminals 291, remote laptops 292, remote desktops 293, and remote web servers 294.

FIG. 2 illustrates an exemplary network environment. Those of ordinary skill in the art will appreciate that the teachings of the present invention can be used with any number of network environments and network configurations.

The Present Invention

The present invention teaches methods and systems to improve computer forensics with search and machine learning. This invention allows organizations that anticipate the need for forensic analysis to prepare in advance by keeping small amounts of information about any content on computer systems, such as files, database entries or schema, or network traffic, as the content is created, deleted, modified, copied or transmitted or received. Computational and storage costs are expanded in advance, which allows faster, better and less expensive computer forensics investigations.

The present invention provides a novel proactive approach for computer forensic investigations. For any type of content that is created, deleted, modified, copied, transmitted or received, a small amount of information about the content, called a signature, is created and stored away. A signature contains one or more fingerprints and other information associated with the target. A fingerprint is a relatively small number of bits, as compared to the size of the file that is computed based on the content of a target. The target can be any file, any file that is owned by a user, any operating system file, any file that is part of a proprietary information system, any file that is related to a network intrusion attack, any database entry or definition, or network traffic. For a text file, for example, a signature contains one or more fingerprints computed based on the content of the file along with other information associated with the file, such as the file name, date and time of record, user/owner information, and fingerprint history. For a database entry or definition, the signature contains one or more fingerprints that are calculated based on the content of the database entry or definition along with other information associated with the database entry or definition. For network traffic, the signature contains one or more fingerprints that are calculated based on the content of the network traffic and along with other information associated with the network traffic, such as time and date information, sender and recipient network addresses, and network protocol.

The fingerprints of the present invention are digital digests of the content of a target. In the hash method, all bits that make up a file are considered as the content of a file. In the present invention, however, the content of a target is defined and represented by selections of tokens that are logically selected from the target. As an example, the content of a target that contains textual information can be defined by a selection of words and phrases within the target. For targets that lack a semantic meaning, idiosyncratic characteristics of the target can be identified and used to represent the contents. Fingerprints are small, taking up a small amount of storage space, when compared to the original content of the target. Fingerprints are also easy to compute, and can identify a file, a database entry or definition, or network traffic by its content as defined by the list of selected tokens. Fingerprints can accommodate small modifications of the file (e.g., small edits or reformatting of a file may not alter its fingerprint). The fingerprints of a minimally edited version of a file mostly or fully match the fingerprints of the original file.

The creation of a signature usually comprises four steps. First, a set of tokens of interest are extracted from a target, such as a file, database entry or definition, or network traffic. Second, the token set undergoes a predetermined sequence of

processing, such as sorting and filtering. Third, a fingerprint is then generated for each retained token set. Lastly, the fingerprint is combined with other information associated with the target file, database entry or definition, or network traffic to generate a signature.

Using a document that contains text information as an example, the first step involves parsing the document, extracting text information and retaining tokens of interest. Tokens of interest may include, but are not limited to, all words, phrases, selective parts of speech, e.g., nouns (names, places, etc.), words longer than a fixed number of characters, words not found in a dictionary, words found within a certain set of predefined lists of words, words of a “foreign nature”, words based on inverse document frequencies (histograms), in other words, words based on collection statistics, and acronyms.

Processing the token set may involve sorting the token set, and may further include filtering the token set. Sorting the token set can be based on, but not limited to, Unicode (alphabetical) ordering, biased weighting on inverse document frequency, and phrase or word length. Filtering the token set and retaining a subset of the tokens can be based on, but not limited to, rules such as selecting the top X % of the tokens, (i.e., $X \geq T1$); or middle tokens, (i.e., $T2 \geq X \geq T1$); or bottom tokens, i.e., $X \leq T2$); or selective sets of tokens, (i.e., every t tokens, e.g., third, seventh, etc.); or no filtration at all, namely retaining all tokens. The retained tokens may be sorted again as previously described. However, sorting is unnecessary if one wishes to retain the same sorting conditions as used previously.

Creating one or more fingerprints of the retained token list can follow several computational methods. For example, a hash based method, where using a hash function, one can encode the sorted list of retained tokens and generate a unique hash for the retained token list. Many popular hash functions can be used for the calculation of the hash, such as MD5, SHA-1, RIPMED, WIRLPOOL, and the variations of these hash functions. Using a hash method for fingerprint creation is advantageous as it calculates quickly, and saves space. However, hash methods are not reversible (i.e., given a hash code, it is computationally impractical to retrieve the original token list).

Another method for fingerprint creation is a bit vector method, which uses a bit vector to encode the presence or absence of retained tokens. The bit vector could be a binary vector using a sequence of Boolean values, each stored as a single bit, or a non-binary numeric vector. The advantage of the bit vector method is that it is a reversible process, but bit vectors are often more costly in terms of storage space.

The creation of fingerprints is generally some form of lossy compression based on a predefined operation. However, it is within the scope of the invention to use a lossless compression method. For multimedia content, such as an image, sound and/or video file, mathematical transformations can be used to create fingerprints. It is apparent to those skilled in the art that fingerprint creation can be achieved through a variety of methods, and are not limited to the above mentioned approaches. Once the fingerprints are created, other information associated with the document is extracted and combined with the fingerprints to create a signature. The other associated information may be information about the document that is accessible through the operating system, which may include, but is not limited to, file name, date and time of record, user/owner information, access history, and fingerprint history. Other information may also include information about the document accessible through an application, which may include, but not limited to, author, time of editing, number words, title, subject, comments, and any other customiz-

able fields or application specific information. There are numerous possibilities regarding the information that can be incorporated into a signature. A person skilled in the art could choose to incorporate any number of desired attributes of the target into a signature, depending on the specific implementation.

FIG. 3 shows an exemplary diagram of the process of generating a signature for a document. The document is first parsed and non-textual information is removed. A set of tokens **311** are extracted **310** from the document. One ordinarily skilled in the art would appreciate that there are a number of other acceptable ways to perform the extraction of the token list. The token set is then processed to yield a unique token list. In the FIG. 3 embodiment, the processing of the token set involves sorting the token set **320**, which produces a sorted list of tokens **323**, and filtering the token set **324**, which generates one or more filtered lists of tokens **325**. One ordinarily skilled in the art would appreciate that there are a number of other acceptable ways to perform the processing of the token set. The retained tokens are then used to generate one or more fingerprints of the document **330**. In one embodiment of the invention, a hash or bit vector can be calculated for the entire list of retained tokens and used as a fingerprint. In another embodiment of the invention, the processed token list can be presented in the form of several subsets of tokens. A hash or bit vector can be calculated for each of the subset of tokens, and the document is represented with a list of fingerprints corresponding to each retained subset of tokens. In yet another embodiment of the invention, a hash or bit vector is calculated for each retained token, and the document is represented with a list of fingerprints corresponding to each retained token. A signature is created **340** by combining other information associated with the document **331** with one or more fingerprints. The resulting signature is then stored.

When a document is modified, if the modification is small, the fingerprint of the file might not change, and the signature is updated with relevant other information. If modifications to a document are not small, then the modified document’s fingerprint may not be sufficiently close to the original fingerprint. After such modification, a new candidate fingerprint is created and compared to the original fingerprint. If sufficient change has occurred in the document, and the candidate fingerprint does not match the original fingerprint, the new candidate fingerprint is added to the document’s signature. The signature may encode other information, including but not limited to information related to derivation. In other embodiments of the invention, similarity may be measured by comparing fingerprints, signatures or both.

FIG. 4 illustrates document modification and further fingerprint generation. When a document is modified **410**, a new candidate fingerprint is generated **420** based on the content of the modified document using the method exemplified in FIG. 3. The new candidate fingerprint is then compared with the fingerprint representing the original version of the document **430**. The actual original document does not need to be retrieved for comparison. If the candidate fingerprint does not differ from the original document, the modification of the document is minor. The original fingerprint is then combined with updated other information associated with the document **450** and the updated signature is stored. If the candidate fingerprint differs from the original document, a major modification has occurred. The candidate fingerprint is then added to the original fingerprint **440**. A new signature of the modified document is then created, incorporating the updated other information of the document and stored. If a fingerprint history is implemented in the signature, it is also updated.

The present invention can easily be adapted to other types of files. It is also possible to construct similar fingerprints for multimedia files such as image, video, and sound files. A variety of mathematical transformations can be used to create fingerprints for these file types, such as Laplace transform, Karhunen-Loeve transform and Fourier transform. Metadata text of sound, image, and video can be used to generate fingerprints. Close captioning within a video file is text which can be used to generate fingerprints, as previously described. Speech can be converted to text using existing software tools. Text thus derived can be used to generate fingerprints. Moreover, the digital content of these files can be encoded as a sequence of tokens, like text documents. Executables and dynamically linked libraries (DLL) can be represented as a sequence of tokens, which can be used to produce fingerprints. Text embedded in these files can also be used to create fingerprints. Reverse engineered programs (e.g., Java) can be treated as text. Byte-code languages and scripting languages (e.g., Perl, python) can also be treated as text. The fingerprint creation process produces a relatively small amount of bits, as compared to the original file, and serves as a digest of the content of the original file. A person skilled in the art will appreciate that numerous methods can be used for achieving fingerprint creation. The fingerprint creation process in general is a lossy compression process. However, lossless compression schemes can also be adopted for the fingerprint creation process.

The signatures are stored in a manner preventing a regular user from modifying or deleting the signatures. Because the signatures are used for forensic purposes, their generation and storage is preferably transparent to the regular user. Only authorized personnel and forensic investigators can have access to the stored signatures. In a network environment, signatures can be created on a user system and offloaded to a network server for storage. Signatures can also be stored on a local file system, while denying user access through use of hidden files or hidden partitions. The signatures can also be embedded in encrypted files. One can also use write-once, read-many media for storing signatures. Only authorized personnel or forensic investigators can recover the storage media and be responsible for safe keeping. Off site storage of the signatures may also be desirable. Cryptographic logging mechanisms can be implemented to control and monitor the access of the signatures.

The present invention can be implemented in a variety of ways. In a stand alone system, such as an individual PC, laptop, mobile device (e.g., cell phone, PDA, etc.), signature information is stored locally. In a system that has access to shared file systems, such as file servers, database servers, and network attached storage (NAS), signature information is stored locally or on the shared file systems. In a network based implementation, any system with a network connection can have signature information stored on remote servers. One skilled in the art will appreciate that signatures can be stored in a variety of ways depending on the system or the network configurations of a particular environment.

Fingerprints can be created for information that is stored in any database and also database definitions. Signatures for each database entry are based on content and can be created for the entire database. As an example, signatures can be created for emails stored within a server database, allowing the tracing of email senders and receivers. Database definitions, such as schema, relations, tables, keys, and data domains can also have signatures created. When a data manipulation or definition event occurs, such as create table, drop table, or alter table, a new signature is created and stored.

In addition to files, signatures can be created for other applications. Changes to virtual machine file systems could be indexed as changes occur. Contents of removable media could have signatures created during mounting or un-mounting (during connection and disconnection) to a computer system. Compressed or archived files could be parsed and have signatures created.

It can also be useful to create and store signatures for network traffic. For example, signatures can be created for emails entering and exiting a network. Email attachments can have separate signatures created. Network traffic can thus be linked to particular emails and files when stored. Contents of instant message conversations and contents of file transfers can also be used to create signatures for the particular network activity. Signatures can also be created for text messages such as the ones based on Short Message Service (SMS) protocol. Web pages can also have signatures generated. When integrated over time, a digest or profile of one or more user's Internet browsing history can be generated. A person skilled in the art will appreciate that any information or signal transmitting protocol can be used as a target for signature creation. In one embodiment of the invention, a proxy firewall is used, and signatures are created of network traffic passing through. Network policies can be configured so that the network traffic passing through the proxy firewall is not encrypted. When so configured, secure connections are established between an inside user computer to the proxy firewall, and the proxy firewall to an outside server using an encryption protocol such as Transport Layer Security (TLS) or Secure Sockets Layer (SSL). Network traffic encryption only occurs between the inside user computer and the proxy firewall, and between the proxy firewall and the outside server. Contents passing through the proxy firewall are not encrypted and can, therefore, have signatures created. Signatures are stored among other information associated information regarding the network traffic with IP addresses used in communication, therefore facilitating the identification of the origin and destination of the traffic.

Once signatures are stored, there are a variety of methods to analyze them. Similarity between signatures can be ascertained by comparing the signature or the fingerprints for exact matching, percentage of matching, probability of matching, or other mathematical calculation revealing the divergence of the signatures or fingerprints. In one embodiment of the invention, a latent analysis can be performed. Particular signatures and/or fingerprints on individual machines locally or remotely can be searched and compared. Signatures or fingerprints that are stored in a database can be similarly searched. In another embodiment of the invention, an active analysis is performed. Instead of simply searching with signatures and fingerprints, advance or retrospective analysis of the signatures and fingerprints can be performed for the purpose of data mining, user profiling, trend analysis, and anomaly detection.

FIG. 5 presents an exemplary method for performing a latent search. When provided with a signature of interest, the signature can then be used directly as a query signature. Where a document of interest is provided, a query signature can be created **510** using the method exemplified in FIG. 3. Stored signatures are then retrieved from storage **520** and compared to the query signature **530**. The comparison can be performed on signatures, the fingerprints within the signatures, or both. Similarity of the query signature to any stored signature is then determined. If the fingerprints are calculated using a hash method, the similarity is estimated based on hash matches. If the fingerprints are calculated using a bit vector method, the similarity is estimated based on bit vector corre-

lation. If the comparison identifies any stored signatures having similarity above a predetermined threshold, the similar signatures are output for further processing **540**. Other information within the stored signatures similar to the query signature is extracted **550**. Other documents containing content similar to the document of interest, computer systems housing the document of interest or any similar documents, and users that had possession of the document of interest or any similar documents, can all be identified **560**.

FIG. **6** presents an exemplary method for user misuse detection. When a user performs an operation to a document that is within a list of predetermined operations, such as create, modify, copy, move, or delete a document, the system captures this user operation **610**, and a new signature is created **620** and stored **630**. This new signature is then used as a query signature, and compared with stored signatures **640**. In one embodiment of the invention, a subset of all stored signatures, such as signatures of known documents containing classified or sensitive information, or illegal content can be used. If the comparison does not identify any stored signature within this subset having similarity to the query signature above a certain threshold, the user is presumably not manipulating classified, sensitive, or illegal content. No action needs to be taken, the operation proceeds as normal. If the comparison identifies any stored signature within this subset that has similarity to the query signature above a certain threshold, the user is presumed to be manipulating classified, sensitive, or illegal content. A further inquiry whether the user is expected to manipulate such content is performed **650** based on criteria such as security clearance, job assignment, or special permission. If the user is determined to have proper access permission, and is expected to manipulate such content, the operation proceeds as normal. However, if the user does not have proper permission, or is not expected to manipulate such content, then the suspect content is identified based on the query and the stored similar fingerprint or signature **660**, and a misuse alert is sent to authorized personnel or a forensic investigator **670**.

FIG. **7** presents another exemplary method for user misuse detection. All the files that belong to or are accessed by a user are identified based on ownership information and access information **710**. Signatures of the entire collection of these files can be used to generate a user profile for the user **720** and are stored **730**. An updated user profile is then generated at a later time, either by request or based on a periodic schedule. The newly generated user profile is then compared to any or all of the stored user profiles of the same user at earlier times **740**. If no difference above a certain threshold is detected among the user profiles, there is no deviation in user behavior. However, if the newly generated user profile differs from the stored user profile above a certain threshold, a further inquiry is performed to determine whether there is a legitimate reason for such deviation of user behavior **750**. If a legitimate reason is found, such as change in job assignment or upgrade of security clearance, the operation proceeds as normal. If no legitimate reason is found for the deviation of user behavior, the content of the mismatched signatures is identified **760**, and an alert of possible user misuse is sent to authorized personnel or to a forensic investigator **770**.

FIG. **8** presents an exemplary method for detection of unauthorized network communication of sensitive information. When a network server receives inbound or outbound network traffic **810**, a signature is then calculated based on the content of the network traffic **820** and stored **830**. The signature is then used as a query signature and is compared to any previously stored signatures **840**. In one embodiment of the invention, if the query signature has similarity to any stored

signature above a certain threshold, it is then compared to a subset of all stored signatures, such as signatures of known documents containing classified or sensitive information, or illegal content **850**. If the query signature does not have similarity above a certain threshold to any of the subsets of stored signatures, no classified, sensitive, or illegal content is detected. Network traffic is allowed to proceed as normal **860**. However, if classified, sensitive, or illegal content is detected, suspect content and user information is identified **870**, the network traffic is then quarantined **880**, and an alert is sent to an authorized personnel or to a forensic investigator **890**.

This proactive approach makes investigations faster, easier, and less expensive. Given one document, all systems containing that or similar documents can be found quickly and easily. This is true even if the given document is a hard copy. Text information can be extracted from the hard copy either automatically (e.g., scanned, segmented, and converted to text using optical character recognition) or manually (e.g., transcribed by hand into a computer readable format) and used to create a query signature. The present invention can identify systems where a document once existed, even if it is now or otherwise deleted. In classified computer networks not connected to the Internet, such as those employed by government intelligence agencies and defense contractors, strict control of content entering and leaving the classified network is necessary. However, traditionally, there generally is no effective mechanism to track the flow of information within the classified network. The present invention can locate any content within the classified network, and provide a system-wide tracking of any content of interest. In one embodiment of the invention, a real time, system-wide map of the distribution of any particular content can be generated and monitored.

This invention can also be used for evidence discovery. Given one user or a set of users, a forensic analysis could determine documents of interest. Those identified documents could be used to seed a fingerprint search across all systems. That would rapidly identify which other systems needed further consideration for analysis. The present invention can determine the source of files that were not permanently stored, such as temporary files deleted without a user's knowledge.

This invention can be further used for misuse detection. Many systems log accesses to restricted material. However, restricted material is usually defined by its location within the file system, or by other attributes of the file. Once the restricted material leaves the protected file systems location, or loses its original attributes, access logging will no longer be able to detect misuse of the restricted material. The present invention, however, can detect when the access logging fails by verifying that documents that should have been logged were logged. Collection statistics and fingerprints can determine when a document is atypical for a user, which may be a sign of document misuse. The present invention can also help to determine the source of leaks by identifying the systems within which a leaked document was present, and a time line that tracks movement the leaked document through a network.

This invention can also be used for intrusion response. When an intrusion is discovered, the signatures of files associated with the intrusion can be recovered. Even if the original files are deleted, the signatures can still be recovered based on time stamps. These recovered signatures can be used to examine across systems for similar intrusions, and also provide early detection to prevent intrusion from similar attacks.

FIG. **9** illustrates an exemplary system of the present invention. The system of FIG. **9** comprises four components: 1) a

processor for creating/generating fingerprints and signatures for a target, such as a document 910; 2) an extension module to the operating system (OS) configured to trigger signature generation upon occurrence of a certain action 920; 3) a mechanism for storing the generated signatures 930; and 4) a comparator for querying the system for stored signatures and comparing those retrieved for similarity 940. The implemented system may have either a Windows service or Linux daemon running in the background to generate a signature for a given file and to store it, and then to query the stored signatures to determine similarity with other signatures. The system runs with administrator or root privileges.

The extension module of the operating system has several components. First the configuration information must be stored on the system. In Windows, this would be registry entries or configuration files. In Linux, a configuration file is used, which is stored in/etc or another location. The configuration information includes mechanisms for signature creation, other information to store with signatures, mechanism and location for signature storage, events that trigger signature creation and mechanisms for extracting text based on file type. Separate programs or modules can be called to perform text extraction. In Windows, the COM model can be used to extract text from Office documents. In Linux, various utilities can be used to extract text from different file types.

The signature creation is linked into the OS so that signatures are created when desired system events occur, such as file deletion, file copy between file systems, and file modification. As soon as the computer system starts, certain system events are remapped to invoke the signature creation process, and the system waits for the occurrence of these events. When any one of these events is captured, the OS invokes calls to the signature creation process. In Linux, this can be achieved by a loadable kernel module. In Windows, this can be done through a variety of ways. When called, the system identifies the digital object (file) that triggered the operation, and passes a copy or pointer to the file for processing to the fingerprints creation process. Tokens are extracted from the file and processed, fingerprints are generated for the retained token list, other information associated with the file (metadata) is incorporated with the fingerprints, and a signature is generated, all based on the criteria specified in the configuration information.

A basic system can incorporate the entire index of retained tokens (i.e., without filtration). In this embodiment of the invention, a simple tokenization of a document may include converting the entire document to lower-case (remove case sensitive information) and obtaining individual tokens. A token for this basic system is any string of length-4 or more separated by either white space or any form of punctuation. The individual tokens are then sorted according to Unicode ordering to obtain unique tokens. A hash code or bit vector is then generated for each token in the sorted unique token list. In another embodiment of the invention, the same process is used for tokenization of a document and sorting of the unique token list. The process also includes the filtering of the unique token list. Subsets of the unique token list are created based on a list of criteria including, but not limited to, keeping tokens of only 6 characters or longer in length, keeping tokens numbered (in order) 25-50, keeping every 7th token, keep every 25th token, or other similar rules. A hash code or bit vector is then generated for each subset of tokens.

Fingerprints may vary in complexity. A signature created based on a complete index of retained tokens, such as a list sorted according to Unicode, can be highly precise but support only minimal variance. The precision and tolerance to variance of a signature created based on a filtered index of

retained tokens depends on the degree of filtration. A signature based on a highly filtered index provides high recall but low precision. The number of filters employed to generate signatures also affects the complexity. Multiple filters increase precision but also increase the time required for signature calculation and the storage space needed for signature safe-keeping.

A mechanism for storing signatures should be resilient against modification by users. Once the signature is created, it is stored securely. A user other than authorized personnel or a forensic investigator should have no means to modify or delete any signature entry. The signatures can be inserted into a database, allowing for easy queries and off-system storage. Alternatively, signatures can be stored in flat files having only root or administrator permissions.

When given a signature, one can check to see if the signature is in the store. If given a file or document, text is extracted from the file, fingerprints are created, then a signature, and the created query signature is checked against the store. If multiple fingerprints are used to represent a file, any or all of the fingerprints can be used to determine similarity above a predetermined threshold. A proper or predetermined threshold can be the matching of all or some of the fingerprints, a probabilistic analysis of the similarity of the fingerprints, or any other mathematical analysis directed to signature divergence. The higher the threshold, the lower the rate of false positives; however, the higher the rate of false negatives.

These and other advantages of the present invention will be apparent to those skilled in the art from the foregoing specification. Accordingly, it will be recognized by those skilled in the art that changes or modifications may be made to the above-described embodiments without departing from the broad inventive concepts of the invention. It should therefore be understood that this invention is not limited to the particular embodiments described herein, but is intended to include all changes and modifications that are within the scope and spirit of the invention.

What is claimed:

1. A computerized method of proactively generating and comparing computer forensic evidence for a computer system, comprising the steps of:

generating at least one signature for at least one target based on the content of the target; and

comparing the at least one generated signature with at least one previously generated signature for the same at least one target to determine whether the signatures have similarities above a predetermined threshold,

wherein both the at least one generated signature and the at least one previously generated signature are proactively generated for computer forensic evidence for the at least one target and configured to allow a forensic analysis with the computer forensic evidence, and

wherein generating the at least one signature further includes the step of extracting a set of tokens representing a subset of the content of the at least one target.

2. The method of claim 1, wherein the target is a file.

3. The method of claim 2, wherein the file is owned by a user.

4. The method of claim 2, wherein the file is related to a network intrusion attack.

5. The method of claim 2, wherein the at least one signature is generated upon occurrence of a predetermined operation, and the predetermined operation is one or more of creating, deleting, renaming, editing, moving, updating, linking, merging, modifying and copying the file.

6. The method of claim 1, wherein the target is a database entry.

19

7. The method of claim 6, wherein the at least one signature is generated upon occurrence of a predetermined operation, and the predetermined operation is one or more of select, insert, update, delete, merge, begin work, commit, rollback, create, drop, truncate, and alter the database entry.

8. The method of claim 1, wherein the target is a database definition.

9. The method of claim 8, wherein the at least one signature is generated upon occurrence of a predetermined operation, and the predetermined operation is one or more of create, drop, and alter the database definition.

10. The method of claim 1, wherein the target is network traffic.

11. The method of claim 10, wherein the at least one signature is generated upon occurrence of a predetermined operation, and the predetermined operation is when the network traffic enters a network or when the network traffic leaves a network.

12. The method of claim 10, wherein the network traffic may be any one or more of a signal protocol, an email, an attachment of an email, an instant message conversation, a text message, a remote login, a virtual private network, a viewed webpage, a file transfer and file sharing of the network traffic.

13. The method of claim 1, wherein generating the at least one signature for the at least one target comprises the steps of: processing the set of tokens; generating a fingerprint from the set of tokens; and generating the signature for the at least one target by combining the fingerprints with other related information of the at least one target.

14. The method of claim 13, wherein processing the set of tokens includes sorting the set of tokens.

15. The method of claim 14, wherein processing the set of tokens further includes filtering the set of tokens.

16. The method of claim 13, wherein generating the fingerprint from the set of tokens involves a hash method.

17. The method of claim 13, wherein generating the fingerprint from the set of tokens involves implementation of a bit vector method.

18. The method of claim 13, wherein the other related information of the target is accessible by an operating system, and is at least one of file name, date of record, time of record, user or owner information, network address, network protocol, access history and fingerprint history of the target.

19. The method of claim 13, wherein the other related information of the target is accessible by an application.

20. The method of claim 1, further comprising the step of storing the at least one generated signature in a manner preventing deletion or modification of the at least one generated signature by a user other than authorized personnel or a forensic investigator.

21. The method of claim 1, wherein the at least one generated signature is available to authorized personnel or a forensic investigator with access, rights.

22. The method of claim 1, where in the at least one generated signature and respective target are stored on the same computer system.

23. The method of claim 1, where in the at least one generated signature is stored on a first computer system and the at least one target is stored on a second computer system accessible through a computer network.

24. The method of claim 13, further comprising a step of generating a signature for a media file.

20

25. The method of claim 24, wherein the media file comprises a video file.

26. The method of claim 25, wherein the signature is generated based on meta data of the video file.

27. The method of claim 26, wherein the signature for the video file is generated based on length information included in the meta data.

28. The method of claim 25, wherein the signature for the video file is generated based on a closed caption of the video file.

29. The method of claim 25, wherein the signature for the video file includes frames corresponding to feature points of the video file.

30. The method of claim 25, wherein the signature for the video file includes frames corresponding to scene changes of the video file.

31. The method of claim 24, wherein the media file comprises an audio file.

32. The method of claim 31, wherein the audio file includes an audio file selected from the group consisting essentially of: a music file and a speech file.

33. The method of claim 31, wherein the step of generation a signature for the audio file includes a step of generating a transcript based on the audio file and a step of generating the signature for the audio file based on the transcript.

34. A non-transitory computer storage medium that configures a computer system to perform a method of proactively generating and comparing computer forensic evidence for a computer system, the method comprising the steps of:

generating at least one signature for at least one target based on the content of the target; and

comparing the at least one generated signature with at least one previously generated signature for the same at least one target to determine whether the signatures have similarities above a predetermined threshold,

wherein both the at least one generated signature and the at least one previously generated signature are proactively generated for computer forensic evidence for the at least one target and configured to allow a forensic analysis with the computer forensic evidence, and

wherein generating the at least one signature further includes the step of extracting a set of tokens representing a subset of the content of the at least one target.

35. Apparatus for proactively generating and comparing computer forensic evidence; comprising:

a processor, CPU and memory arranged to generate at least one signature for at least one target based on the content of the target; and

a comparator configured to compare the at least one generated signature with at least one previously generated signature for the same at least one target to determine whether the signatures have similarities above a predetermined threshold,

wherein both the at least one generated signature and the at least one previously generated signature are proactively generated for computer forensic evidence for the at least one target and configured to allow a forensic analysis with the computer forensic evidence, and

wherein the processor extracts a set of tokens representing a subset of the content of the at least one target to generate the at least one signature.